arXiv:1411.1152v3 [q-fin.EC] 7 May 2016

Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models*

Ignacio Esponda
(WUSTL)

Demian Pouzo
(UC Berkeley)

May 10, 2016

Abstract

We develop an equilibrium framework that relaxes the standard assumption that people have a correctly-specified view of their environment. Each player is characterized by a (possibly misspecified) subjective model, which describes the set of feasible beliefs over payoff-relevant consequences as a function of actions. We introduce the notion of a Berk-Nash equilibrium: Each player follows a strategy that is optimal given her belief, and her belief is restricted to be the best fit among the set of beliefs she considers possible. The notion of best fit is formalized in terms of minimizing the Kullback-Leibler divergence, which is endogenous and depends on the equilibrium strategy profile. Standard solution concepts such as Nash equilibrium and self-confirming equilibrium constitute special cases where players have correctly-specified models. We provide a learning foundation for Berk-Nash equilibrium by extending and combining results from the statistics literature on misspecified learning and the economics literature on learning in games.

*We thank Vladimir Asriyan, Pierpaolo Battigalli, Larry Blume, Aaron Bodoh-Creed, Sylvain Chassang, Emilio Espino, Erik Eyster, Drew Fudenberg, Yuriy Gorodnichenko, Stephan Lauermann, Natalia Lazzati, Kristóf Madarász, Matthew Rabin, Ariel Rubinstein, Joel Sobel, Jörg Stoye, several seminar participants, and especially a co-editor and four anonymous referees for very helpful comments. Esponda: Olin Business School, Washington University in St. Louis, 1 Brookings Drive, Campus Box 1133, St. Louis, MO 63130, iesponda@wustl.edu; Pouzo: Department of Economics, UC Berkeley, 530-1 Evans Hall #3880, Berkeley, CA 94720, dpouzo@econ.berkeley.edu.

1 Introduction

Most economists recognize that the simplifying assumptions underlying their models are often wrong. But, despite recognizing that models are likely to be misspecified, the standard approach (with exceptions noted below) assumes that economic agents have a correctly specified view of their environment. We present an equilibrium framework that relaxes this standard assumption and allows the modeler to postulate that economic agents have a subjective and possibly incorrect view of their world.

An objective game represents the true environment faced by the agent (or players, in the case of several interacting agents). Payoff relevant states and privately observed signals are drawn from an objective probability distribution. Each player observes her own private signal and then players simultaneously choose actions. The action profile and the realized state determine consequences, and consequences determine payoffs.

In addition, each player has a subjective model representing her own view of the environment. Formally, a subjective model is a set of probability distributions over own consequences as a function of a player's own action and information. Crucially, we allow the subjective model of one or more players to be misspecified, which roughly means that the set of subjective distributions does not include the true, objective distribution. For example, a consumer might perceive a nonlinear price schedule to be linear and, therefore, respond to average, not marginal, prices. Or traders might not realize that the value of trade is partly determined by the terms of trade.

A *Berk-Nash equilibrium* is a strategy profile such that, for each player, there exists a belief with support in her subjective model satisfying two conditions. First, the strategy is optimal given the belief. Second, the belief puts probability one on the set of subjective distributions over consequences that are "closest" to the true distribution, where the true distribution is determined by the objective game and the actual strategy profile. The notion of "closest" is given by a weighted version of the Kullback-Leibler divergence, also known as relative entropy.

Berk-Nash equilibrium includes standard and boundedly rational solution concepts in a common framework, such as Nash, self-confirming (e.g., Battigalli (1987), Fudenberg and Levine (1993a), Dekel et al. (2004)), fully cursed (Eyster and Rabin, 2005), and analogy-based expectation equilibrium (Jehiel (2005), Jehiel and Koessler (2008)). For example, suppose that the game is *correctly specified* (i.e., the support of each player's prior contains the true distribution) and that the game is *strongly iden-*

tified (i.e., there is a unique distribution—whether or not correct—that matches the observed data). Then Berk-Nash equilibrium is equivalent to Nash equilibrium. If the strong identification assumption is dropped, then Berk-Nash is a self-confirming equilibrium. In addition to unifying previous work, our framework provides a systematic approach for extending previous cases and exploring new types of misspecifications.

We provide a foundation for Berk-Nash equilibrium (and the use of Kullback-Leibler divergence as a measure of "distance") by studying a dynamic setup with a fixed number of players playing the objective game repeatedly. Each player believes that the environment is stationary and starts with a prior over her subjective model. In each period, players use the observed consequences to update their beliefs according to Bayes' rule. The main objective is to characterize limiting behavior when players behave optimally but learn with a possibly misspecified subjective model.[1]

The main result is that, if players' behavior converges, then it converges to a Berk-Nash equilibrium. A converse result, showing that we can converge to any Berk-Nash equilibrium of the game for some initial (non-doctrinaire) prior, does not hold. But we obtain a positive convergence result by relaxing the assumption that players exactly optimize. For any given Berk-Nash equilibrium, we show that convergence to that equilibrium occurs if agents are myopic and make *asymptotically optimal* choices (i.e., optimization mistakes vanish with time).

There is a longstanding interest in studying the behavior of agents who hold misspecified views of the world. Examples come from diverse fields including industrial organization, mechanism design, information economics, macroeconomics, and psychology and economics (e.g., Arrow and Green (1973), Kirman (1975), Sobel (1984), Kagel and Levin (1986), Nyarko (1991), Sargent (1999), Rabin (2002)), although there is often no explicit reference to misspecified learning. Most of the literature, however, focuses on particular settings, and there has been little progress in developing a unified framework. Our treatment unifies both "rational" and "boundedly rational" approaches, thus emphasizing that modeling the behavior of misspecified players does not constitute a large departure from the standard framework.

Arrow and Green (1973) provide a general treatment and make a distinction between objective and subjective games. Their framework, though, is more restrictive than ours in terms of the types of misspecifications that players are allowed to have.

[1]In the case of multiple agents, the environment need not be stationary, and so we are ignoring repeated game considerations where players take into account how their actions affect others' future play. We discuss the extension to a population model with a continuum of agents in Section 5.

Moreover, they do not establish existence or provide a learning foundation for equilibrium. Recently, Spiegler (2014) introduced a framework that uses Bayesian networks to analyze decision making under imperfect understanding of correlation structures.[2]

Our paper is also related to the bandit (e.g., Rothschild (1974), McLennan (1984), Easley and Kiefer (1988)) and self-confirming equilibrium (SCE) literatures, which highlight that agents might optimally end up with incorrect beliefs if experimentation is costly.[3] We also allow beliefs to be incorrect due to insufficient feedback, but our main contribution is to allow for misspecified learning. When players have misspecified models, beliefs may be incorrect and endogenously depend on own actions even if there is persistent experimentation; thus, an equilibrium framework is needed to characterize steady-state behavior even in single-agent settings.[4]

From a technical perspective, we extend and combine results from two literatures. First, the idea that equilibrium is a result of a learning process comes from the literature on learning in games. This literature studies explicit learning models to justify Nash and SCE (e.g., Fudenberg and Kreps (1988), Fudenberg and Kreps (1993), Fudenberg and Kreps (1995), Fudenberg and Levine (1993b), Kalai and Lehrer (1993)).[5] We extend this literature by allowing players to learn with models of the world that are misspecified even in steady state.

Second, we rely on and contribute to the literature studying the limiting behavior of Bayesian posteriors. The results from this literature have been applied to decision problems with correctly specified agents (e.g., Easley and Kiefer, 1988). In particular, an application of the martingale convergence theorem implies that beliefs converge almost surely under the agent's subjective prior. This result, however, does not guarantee convergence of beliefs according to the true distribution if the agent has a misspecified model and the support of her prior does not include the true distribution. Thus, we take a different route and follow the statistics literature on misspecified learning. This literature characterizes limiting beliefs in terms of the Kullback-Leibler divergence (e.g., Berk (1966), Bunke and Milhaud (1998)).[6] We extend the statistics

[2]Some explanations for why players may have misspecified models include the use of heuristics (Tversky and Kahneman, 1973), complexity (Aragones et al., 2005), the desire to avoid over-fitting the data (Al-Najjar (2009), Al-Najjar and Pai (2013)), and costly attention (Schwartzstein, 2009).

[3]In the macroeconomics literature, the term SCE is sometimes used in a broader sense to include cases where agents have misspecified models (e.g., Sargent, 1999).

[4]Two extensions of SCE are also potentially applicable: restrictions on beliefs based on introspection (e.g., Rubinstein and Wolinsky, 1994), and ambiguity aversion (Battigalli et al., 2012).

[5]See Fudenberg and Levine (1998, 2009) for a survey of this literature.

[6]White (1982) shows that the Kullback-Leibler divergence also characterizes the limiting behavior

literature on misspecified learning to the case where agents are not only passively learning about their environment but are also actively learning by taking actions.

We present the framework and examples in Section 2, discuss the relationship to other solution concepts in Section 3, and provide a learning foundation in Section 4. We discuss assumptions and extension in Section 5.

2 The framework

2.1 The environment

A (simultaneous-move) **game** $\mathcal{G} =< \mathcal{O}, \mathcal{Q} >$ is composed of a (simultaneous-move) objective game $\mathcal{O}$ and a subjective model $\mathcal{Q}$.

Objective game. A (simultaneous-move) **objective game** is a tuple

$$\mathcal{O} = \langle I, \Omega, \mathbb{S}, p, \mathbb{X}, \mathbb{Y}, f, \pi \rangle ,$$

where: I is the set of players; Ω is the set of payoff-relevant states; $\mathbb{S} = \times_{i \in I} \mathbb{S}^i$ is the set of profiles of signals, where $\mathbb{S}^i$ is the set of signals of player i; p is a probability distribution over $\Omega \times \mathbb{S}$, and, for simplicity, it is assumed to have marginals with full support; we use standard notation to denote marginal and conditional distributions, e.g., $p_{\Omega|S^i}(\cdot \mid s^i)$ denotes the conditional distribution over Ω given $S^i = s^i$; $\mathbb{X} = \times_{i \in I} \mathbb{X}^i$ is a set of profiles of actions, where $\mathbb{X}^i$ is the set of actions of player i; $\mathbb{Y} = \times_{i \in I} \mathbb{Y}^i$ is a set of profiles of (observable) consequences, where $\mathbb{Y}^i$ is the set of consequences of player i; $f = (f^i)_{i \in I}$ is a profile of feedback or consequence functions, where $f^i : \mathbb{X} \times \Omega \to \mathbb{Y}^i$ maps outcomes in $\Omega \times \mathbb{X}$ into consequences of player i; and $\pi = (\pi^i)_{i \in I}$, where $\pi^i : \mathbb{X}^i \times \mathbb{Y}^i \to \mathbb{R}$ is the payoff function of player i.[7] For simplicity, we prove the results for the case where all of the above sets are finite.[8]

The timing of the objective game is as follows: First, a state and a profile of signals are drawn according to p. Second, each player privately observes her own signal. Third, players simultaneously choose actions. Finally, each player observes her consequence and obtains a payoff. We implicitly assume that players observe at

of the maximum quasi-likelihood estimator.

[7] The concept of a feedback function is borrowed from the SCE literature. Also, while it is redundant to have π^i depend on x^i, it simplifies the notation in applications.

[8] In the working paper version (Esponda and Pouzo, 2014), we provide technical conditions under which the results extend to nonfinite Ω and $\mathbb{Y}$.

least their own actions and payoffs.[9]

A strategy of player i is a mapping $\sigma^i : \mathbb{S}^i \to \Delta(\mathbb{X}^i)$. The probability that player i chooses action x^i after observing signal s^i is denoted by $\sigma^i(x^i \mid s^i)$. A strategy profile is a vector of strategies $\sigma = (\sigma^i)_{i \in I}$; let Σ denote the space of all strategy profiles.

Fix an objective game. For each strategy profile σ, there is an **objective distribution** over player i's consequences, $Q^i_\sigma : \mathbb{S}^i \times \mathbb{X}^i \to \Delta(\mathbb{Y}^i)$, where

$$Q^i_\sigma(y^i \mid s^i, x^i) = \sum_{\{(\omega, x^{-i}) : f^i(x^i, x^{-i}, \omega) = y^i\}} \sum_{s^{-i}} \prod_{j \neq i} \sigma^j(x^j \mid s^j) p_{\Omega \times S^{-i} \mid S^i}(\omega, s^{-i} \mid s^i), \qquad (1)$$

for all $(s^i, x^i, y^i) \in \mathbb{S}^i \times \mathbb{X}^i \times \mathbb{Y}^i$.[10] The objective distribution represents the true distribution over consequences, conditional on a player's own action and signal, given the objective game and a strategy profile followed by the players.

Subjective model. The subjective model represents the set of distributions over consequences that players consider possible a priori. For a fixed objective game, a **subjective model** is a tuple

$$\mathcal{Q} = \langle \Theta, (Q_\theta)_{\theta \in \Theta} \rangle ,$$

where $\Theta = \times_{i \in I} \Theta^i$ and Θ^i is player i's parameter set; and $Q_\theta = (Q^i_{\theta^i})_{i \in I}$, where $Q^i_{\theta^i} : \mathbb{S}^i \times \mathbb{X}^i \to \Delta(\mathbb{Y}^i)$ is the conditional distribution over player i's consequences parameterized by $\theta^i \in \Theta^i$; we denote the conditional distribution by $Q_{\theta^i}(\cdot \mid s^i, x^i)$.[11]

While the objective game represents the true environment, the subjective model represents the players' perception of their environment. This separation between objective and subjective models is crucial in this paper.

Remark 1. A special case of a subjective model is one where each player understands the objective game being played but is uncertain about the distribution over states, the consequence function, and (in the case of multiple players) the strategies of other players. In this special case, player i's uncertainty about p, f^i, and σ^{-i} can be described by a parametric model p_{θ^i}, $f^i_{\theta^i}$, $\sigma^{-i}_{\theta^i}$, where $\theta^i \in \Theta^i$. A subjective distribution $Q^i_{\theta^i}$ is then derived by replacing p, f^i, and σ^{-i} with p_{θ^i}, $f^i_{\theta^i}$, $\sigma^{-i}_{\theta^i}$ in equation (1).[12] □

[9] See Online Appendix E for the case where players do not observe own payoffs.

[10] As usual, the superscript $-i$ denotes a profile where the i'th component is excluded

[11] For simplicity, we assume that players know the distribution over own signals.

[12] In this case, a player understands that other players mix independently but, due to uncertainty over the parameter θ^i that indexes $\sigma^{-i}_{\theta^i} = (\sigma^j_{\theta^i})_{j \neq i}$, she may have correlated beliefs about her oppo-

By defining $Q^i_{\theta^i}$ as a primitive, we stress two points. First, this object is sufficient to characterize behavior. Second, working with general subjective distributions allows for more general types of misspecifications, where players do not even have to understand the structural elements that determine their payoff relevant consequences.

We maintain the following assumptions about the subjective model.

Assumption 1. For all $i \in I$: (i) Θ^i is a compact subset of an Euclidean space, (ii) $Q^i_{\theta^i}(y^i \mid s^i, x^i)$ is continuous as a function of $\theta^i \in \Theta^i$ for all $(y^i, s^i, x^i) \in \mathbb{Y}^i \times \mathbb{S}^i \times \mathbb{X}^i$, (iii) For all $\theta^i \in \Theta^i$, there exists a sequence $(\theta^i_n)_n$ in Θ^i such that $\lim_{n\to\infty} \theta^i_n = \theta^i$ and such that, for all n, $Q^i_{\theta^i_n}(y^i \mid s^i, x^i) > 0$ for all $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$, $y^i \in f^i(x^i, \mathbb{X}^{-i}, \omega)$, and $\omega \in supp(p_{\Omega|S^i}(\cdot \mid s^i))$.

Conditions (i) and (ii) are the standard conditions used to define a *parametric* model in statistics (e.g., Bickel et al. (1993)). Condition (iii) plays two roles. First, it guarantees that there exists at least one parameter value that attaches positive probability to every feasible observation. In particular, it rules out what can be viewed as a stark misspecification in which every element of the subjective model attaches zero probability to an event that occurs with positive true probability. Second, it imposes a "richness" condition on the subjective model: If a feasible event is deemed impossible by some parameter value, then that parameter value is not isolated in the sense that there are nearby parameter values that consider every feasible event to be possible. In Section 5, we show that equilibrium may fail to exist and steady-state behavior need not be characterized by equilibrium without this assumption.

2.2 Examples

We illustrate the environment by presenting several examples that had previously not been integrated into a common framework.[13] In examples with a single agent, we drop the i subscript from the notation.

Example 2.1. ***Monopolist with unknown demand.*** A monopolist faces demand $y = f(x, \omega) = \phi_0(x) + \omega$, where $x \in \mathbb{X}$ is the price chosen by the monopolist

nents' strategies, as in Fudenberg and Levine (1993a).

[13] Nyarko (1991) studies a special case of Example 2.1 and shows that a steady state does not exist in pure strategies; Sobel (1984) considers a misspecification similar to Example 2.2; Tversky and Kahneman's (1973) story motivates Example 2.3; Sargent (1999, Chapter 7) studies Example 2.4; and Kagel and Levin (1986), Eyster and Rabin (2005), Jehiel and Koessler (2008), and Esponda (2008) study Example 2.5. See Esponda and Pouzo (2014) for additional examples.

and ω is a mean-zero shock with distribution $p \in \Delta(\Omega)$. The monopolist observes sales y, but not the shock. The monopolist does not observe any signal, and so we omit signals from the notation. The monopolist's payoff is $\pi(x, y) = xy$ (i.e., there are no costs). The monopolist's uncertainty about p and f is described by a parametric model f_θ, p_θ, where $y = f_\theta(x, \omega) = a - bx + \omega$ is the demand function, $\theta = (a, b) \in \Theta$ is a parameter vector, and $\omega \sim N(0, 1)$ (i.e., p_θ is a standard normal distribution for all $\theta \in \Theta$). In particular, this example corresponds to the special case discussed in Remark 1, and $Q_\theta(\cdot \mid x)$ is a normal density with mean $a - bx$ and unit variance. □

Example 2.2. ***Nonlinear taxation.*** An agent chooses effort $x \in \mathbb{X}$ at cost $c(x)$ and obtains income $z = x + \omega$, where ω is a zero-mean shock with distribution $p \in \Delta(\Omega)$. The agent pays taxes $t = \tau(z)$, where $\tau(\cdot)$ is a nonlinear tax schedule. The agent does not observe any signal, and so we omit them. The agent observes $y = (z, t)$ and obtains payoff $\pi(x, z, t) = z - t - c(x)$.[14] She understands how effort translates into income but fails to realize that the marginal tax rate depends on income. We compare two models that capture this misspecification. In model A, the agent believes in a random coefficient model, $t = (\theta^A + \varepsilon)z$, in which the marginal and average tax rates are both equal to $\theta^A + \varepsilon$, where $\theta^A \in \Theta^A = \mathbb{R}$. In model B, the agent believes that $t = \theta_1^B + \theta_2^B z + \varepsilon$, where θ_2^B is the constant marginal tax rate and $\theta^B = (\theta_1^B, \theta_2^B) \in \Theta^B = \mathbb{R}^2$.[15] In both models, $\varepsilon \sim N(0, 1)$ measures uncertain aspects of the schedule (e.g., variations in tax rates or credits). Thus, $Q_\theta^j(t, z \mid x) = Q_\theta^j(t \mid z)p(z - x)$, where $Q_\theta^j(\cdot \mid z)$ is a normal density with mean $\theta^A z$ and variance z^2 in model $j = A$ and mean $\theta_1^B + \theta_2^B z$ and unit variance in model $j = B$. □

Example 2.3. ***Regression to the mean.*** An instructor observes the initial performance s of a student and decides to praise or criticize him, $x \in \{C, P\}$. The student then performs again and the instructor observes his final performance, s'. The truth is that performances $y = (s, s')$ are independent, standard normal random variables. The instructor's payoff is $\pi(x, s, s') = s' - c(x, s)$, where $c(x, s) = \kappa |s| > 0$ if either $s > 0, x = C$ or $s < 0, x = P$, and, in all other cases, $c(x, s) = 0$.[16] The function c represents a (reputation) cost from lying (i.e., criticizing above-average

[14] Formally, $f(x, \omega) = (z(x, \omega), t(x, \omega))$, where $z(x, \omega) = x + \omega$ and $t(x, \omega) = \tau(x + \omega)$.

[15] It is not necessary to assume that Θ^A and Θ^B are compact for an equilibrium to exist; the same comment applies to Examples 2.3 and 2.4

[16] Formally, $\omega = (s, s')$, p is the product of standard normal distributions, and $y = f(x, \omega) = \omega$.

performances or praising below-average ones) that increases in the size of the lie. Because the instructor cannot influence performance, it is optimal to praise if $s > 0$ and to criticize if $s < 0$. The instructor, however, does not admit the possibility of regression to the mean and believes that $s' = s + \theta_x + \varepsilon$, where $\varepsilon \sim N(0,1)$, and $\theta = (\theta_C, \theta_P) \in \Theta$ parameterizes her perceived influence on performance.[17] Thus, letting $\bar{Q}_\theta(\cdot \mid x)$ be the a normal density with mean $s + \theta_x$ and unit variance, it follows that $Q_\theta(\hat{s}, s' \mid s, x) = \bar{Q}_\theta(s' \mid s, x)$ if $\hat{s} = s$ and 0 otherwise. □

Example 2.4. ***Monetary policy.*** Two players, the government (G) and the public (P), i.e., $I = \{G, P\}$, choose monetary policy x^G and inflation forecasts x^P, respectively. They do not observe signals, and so we omit them. Inflation, e, and unemployment, U, are determined by[18]

$$e = x^G + \varepsilon_e \tag{2}$$
$$U = u^* - \lambda(e - x^P) + \varepsilon_U, \tag{3}$$

where $u^* > 0$, $\lambda \in (0,1)$ and $\omega = (\varepsilon_e, \varepsilon_U) \in \Omega = \mathbb{R}^2$ are shocks with a full support distribution $p \in \Delta(\Omega)$ and $Var(\varepsilon_e) > 0$. The public and the government observe realized inflation and unemployment, but not the error terms. The government's payoff is $\pi(x^G, e, U) = -(U^2 + e^2)$. For simplicity, we focus on the government's problem and assume that the public has correct beliefs and chooses $x^P = x^G$. The government understands how its policy x^G affects inflation, but does not realize that unemployment is affected by *surprise* inflation:

$$U = \theta_1 - \theta_2 e + \varepsilon_U. \tag{4}$$

The subjective model is parameterized by $\theta = (\theta_1, \theta_2) \in \Theta$, and it follows that $Q_\theta(e, U \mid x^G)$ is the density implied by the equations (2) and (4). □

Example 2.5. ***Trade with adverse selection.*** A buyer with valuation $v \in \mathbb{V}$ and a seller submit a (bid) price $x \in \mathbb{X}$ and an ask price $a \in \mathbb{A}$, respectively. The seller's ask price and the buyer's value are drawn from $p \in \Delta(\mathbb{A} \times \mathbb{V})$, so that $\Omega = \mathbb{A} \times \mathbb{V}$

[17] A model that allows for regression to the mean is $s' = \alpha s + \theta_x + \varepsilon$; in this case, the agent would correctly learn that $\alpha = 0$ and $\theta_x = 0$ for all x. Rabin and Vayanos (2010) study a related setup in which the agent believes that shocks are autoregressive when in fact they are i.i.d.

[18] Formally, $\omega = (\varepsilon_e, \varepsilon_U)$ and $y = (e, U) = f(x^G, x^P, \omega)$ is given by equations (2) and (3).

is the state space. Thus, the buyer is the only decision maker.[19] After submitting a price, the buyer observes $y = \omega = (a, v)$ and gets payoff $\pi(x, a, v) = v - x$ if $a \leq x$ and zero otherwise. In other words, the buyer observes perfect feedback, gets $v - x$ if there is trade, and 0 otherwise. When making an offer, she does not know her value or the seller's ask price. She also does not observe any signals, and so we omit them. Finally, suppose that A and V are correlated but that the buyer believes they are independent. This is captured by letting $Q_\theta = \theta$ and $\Theta = \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$. □

2.3 Definition of equilibrium

Distance to true model. In equilibrium, we will require players' beliefs to put probability one on the set of subjective distributions over consequences that are "closest" to the objective distribution. The following function, which we call the **weighted Kullback-Leibler divergence** (wKLD) function of player i, is a weighted version of the standard Kullback-Leibler divergence in statistics (Kullback and Leibler, 1951). It represents a "distance" between the objective distribution over i's consequences given a strategy profile $\sigma \in \Sigma$ and the distribution as parameterized by $\theta^i \in \Theta^i$:[20]

$$K^i(\sigma, \theta^i) = \sum_{(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i} E_{Q^i_\sigma(\cdot \mid s^i, x^i)} \left[\ln \frac{Q^i_\sigma(Y^i \mid s^i, x^i)}{Q^i_{\theta^i}(Y^i \mid s^i, x^i)} \right] \sigma^i(x^i \mid s^i) p_{S^i}(s^i). \tag{5}$$

The set of closest parameter values of player i given σ is the set

$$\Theta^i(\sigma) \equiv \arg\min_{\theta^i \in \Theta^i} K^i(\sigma, \theta^i).$$

The interpretation is that $\Theta^i(\sigma) \subset \Theta^i$ is the set of parameter values that player i can believe to be possible after observing feedback consistent with strategy profile σ.

Remark 2. We show in Section 4 that wKLD is the right notion of distance in a learning model with Bayesian players. Here, we provide an heuristic argument for a Bayesian agent (we drop i subscripts for clarity) with parameter set $\Theta = \{\theta_1, \theta_2\}$ who observes data over t periods, $(s_\tau, x_\tau, y_\tau)_{\tau=0}^{t-1}$, that comes from repeated play of

[19] The typical story is that there is a population of sellers each of whom follows the weakly dominant strategy of asking for her valuation; thus, the ask price is a function of the seller's valuation and, if buyer and seller valuations are correlated, then the ask price and buyer valuation are also correlated.

[20] The notation E_Q denotes expectation with respect to the probability distribution Q. Also, we use the convention that $-\ln 0 = \infty$ and $0 \ln 0 = 0$.

the objective game under strategy σ. Let $\rho_0 = \mu_0(\theta_2)/\mu_0(\theta_1)$ denote the agent's ratio of priors. Applying Bayes' rule and simple algebra, the posterior probability over θ_1 after t periods is

$$\mu_t(\theta_1) = \left(1 + \rho_0 \Pi_{\tau=0}^{t-1} \frac{Q_{\theta_2}(y_\tau \mid s_\tau, x_\tau)}{Q_{\theta_1}(y_\tau \mid s_\tau, x_\tau)}\right)^{-1} = \left(1 + \rho_0 \Pi_{\tau=0}^{t-1} \frac{Q_{\theta_2}(y_\tau \mid s_\tau, x_\tau)/Q_\sigma(y_\tau \mid s_\tau, x_\tau)}{Q_{\theta_1}(y_\tau \mid s_\tau, x_\tau)/Q_\sigma(y_\tau \mid s_\tau, x_\tau)}\right)^{-1}$$
$$= \left(1 + \rho_0 \exp\left\{-t\left(\frac{1}{t}\sum_{\tau=0}^{t-1} \ln \frac{Q_\sigma(y_\tau \mid s_\tau, x_\tau)}{Q_{\theta_2}(y_\tau \mid s_\tau, x_\tau)} - \frac{1}{t}\sum_{\tau=0}^{t-1} \ln \frac{Q_\sigma(y_\tau \mid s_\tau, x_\tau)}{Q_{\theta_1}(y_\tau \mid s_\tau, x_\tau)}\right)\right\}\right)^{-1}$$

where the second equality follows by multiplying and dividing by $\Pi_{\tau=0}^{t-1} Q_\sigma(y_\tau \mid s_\tau, x_\tau)$. By a law of large numbers argument and the fact that the true joint distribution over (s, x, y) is given by $Q_\sigma(y \mid x, s)\sigma(x \mid s)p_S(s)$, the difference in the log-likelihood ratios converges to $K(\sigma, \theta_2) - K(\sigma, \theta_1)$. Suppose that $K(\sigma, \theta_1) > K(\sigma, \theta_2)$. Then, for sufficiently large t, the posterior belief $\mu_t(\theta_1)$ is approximately equal to $1/(1 + \rho_0 \exp(-t(K(\sigma, \theta_2) - K(\sigma, \theta_1))))$, which converges to 0. Therefore, the posterior eventually assigns zero probability to θ_1. On the other hand, if $K(\sigma, \theta_1) < K(\sigma, \theta_2)$, then the posterior eventually assigns zero probability to θ_2. Thus, the posterior eventually assigns zero probability to parameter values that do not minimize $K(\sigma, \cdot)$. □

Remark 3. Because the wKLD function is weighted by a player's own strategy, it places no restrictions on beliefs about outcomes that only arise following out-of-equilibrium actions (beyond the restrictions imposed by Θ). □

The next result collects some useful properties of the wKLD function.

Lemma 1. *(i) For all $i \in I$, $\theta^i \in \Theta^i$, and $\sigma \in \Sigma$, $K^i(\sigma, \theta^i) \geq 0$, with equality holding if and only if $Q_{\theta^i}(\cdot \mid s^i, x^i) = Q^i_\sigma(\cdot \mid s^i, x^i)$ for all (s^i, x^i) such that $\sigma^i(x^i \mid s^i) > 0$. (ii) For all $i \in I$, $\Theta^i(\cdot)$ is nonempty, upper hemicontinuous, and compact valued.*

Proof. See the Appendix. □

The upper-hemicontinuity of $\Theta^i(\cdot)$ would follow from the Theorem of the Maximum had we assumed $Q^i_{\theta^i}$ to be positive for all feasible events and $\theta^i \in \Theta^i$, since the wKLD function would then be finite and continuous. But this assumption may be strong in some cases.[21] Assumption 1(iii) weakens this assumption by requiring that it holds for a dense subset of Θ, and still guarantees that $\Theta^i(\cdot)$ is upper hemicontinuous.

[21] For example, it rules out cases where a player believes others follow pure strategies.

OPTIMALITY. In equilibrium, we will require each player to choose a strategy that is optimal given her beliefs. A strategy σ^i for player i is **optimal** given $\mu^i \in \Delta(\Theta^i)$ if $\sigma^i(x^i \mid s^i) > 0$ implies that

$$x^i \in \arg\max_{\bar{x}^i \in \mathbb{X}^i} E_{\bar{Q}^i_{\mu^i}(\cdot \mid s^i, \bar{x}^i)} \left[\pi^i(\bar{x}^i, Y^i)\right] \tag{6}$$

where $\bar{Q}^i_{\mu^i}(\cdot \mid s^i, x^i) = \int_{\Theta^i} Q^i_{\theta^i}(\cdot \mid s^i, x^i)\mu^i(d\theta^i)$ is the distribution over consequences of player i, conditional on $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$, induced by μ^i.

DEFINITION OF EQUILIBRIUM. We propose the following solution concept.

Definition 1. A strategy profile σ is a **Berk-Nash equilibrium** of game $\mathcal{G}$ if, for all players $i \in I$, there exists $\mu^i \in \Delta(\Theta^i)$ such that

(i) σ^i is optimal given μ^i, and

(ii) $\mu^i \in \Delta(\Theta^i(\sigma))$, i.e., if $\hat{\theta}^i$ is in the support of μ^i, then

$$\hat{\theta}^i \in \arg\min_{\theta^i \in \Theta^i} K^i(\sigma, \theta^i).$$

Definition 1 places two restrictions on equilibrium behavior: (i) optimization given beliefs, and (ii) endogenous restrictions on beliefs. For comparison, note that the definition of Nash equilibrium is identical to Definition 1 except that condition (ii) is replaced with the condition that players have correct beliefs, i.e., $\bar{Q}^i_{\mu^i} = Q^i_\sigma$.

EXISTENCE OF EQUILIBRIUM. The standard existence proof of Nash equilibrium cannot be used here because the analogous version of a best response correspondence is not necessarily convex valued. To prove existence, we first perturb payoffs and establish that equilibrium exists in the perturbed game. We then consider a sequence of equilibria of perturbed games, where perturbations go to zero, and establish that the limit is a Berk-Nash equilibrium of the (unperturbed) game.[22] The nonstandard part of the proof is to prove existence of equilibrium in the perturbed game. The perturbed best response correspondence is still not necessarily convex valued. Our approach is to characterize equilibrium as a fixed point of a *belief correspondence* and show that it satisfies the requirements of a generalized version of Kakutani's fixed point theorem.

[22]The idea of perturbations and the strategy of the existence proof date back to Harsanyi (1973); Selten (1975) and Kreps and Wilson (1982) also used these ideas to prove existence of perfect and sequential equilibrium, respectively.

Theorem 1. *Every game has at least one Berk-Nash equilibrium.*

Proof. See the Appendix. □

2.4 Examples: Finding a Berk-Nash equilibrium

Example 2.1, continued from pg. 6. *Monopolist with unknown demand.* Let $\sigma = (\sigma_x)_{x\in\mathbb{X}}$ denote a strategy, where σ_x is the probability of choosing price $x \in \mathbb{X}$. Because this is a single-agent problem, the objective distribution does not depend on σ; hence, we denote it by $Q_0(\cdot \mid x)$, which is a normal density with mean $\phi_0(x)$ and unit variance. Similarly, $Q_\theta(\cdot \mid x)$ is a normal density with mean $\phi_\theta(x) = a - bx$ and unit variance. It follows from equation (5) that

$$K(\sigma,\theta) = \sum_{x\in\mathbb{X}} \sigma_x \frac{1}{2} E_{Q_0(\cdot|x)}\left[(Y-\phi_\theta(x))^2 - (Y-\phi_0(x))^2\right] = \sum_{x\in\mathbb{X}} \sigma_x \frac{1}{2}\left(\phi_0(x)-\phi_\theta(x)\right)^2.$$

For concreteness, let $\mathbb{X} = \{2, 10\}$, $\phi_0(2) = 34$ and $\phi_0(10) = 2$, and $\Theta = [33, 40] \times [3, 3.5]$.[23] Let $\theta^0 \in \mathbb{R}^2$ provide a perfect fit for demand, i.e., $\phi_{\theta^0}(x) = \phi_0(x)$ for all $x \in \mathbb{X}$. In this example, $\theta^0 = (a^0, b^0) = (42, 4) \notin \Theta$ and, therefore, we say that the monopolist has a *misspecified* model. The dashed line in Figure 1 depicts optimal behavior: the optimal price is 10 to the left, it is 2 to the right, and the monopolist is indifferent for parameter values on the dashed line.

To solve for equilibrium, we first consider pure strategies. If $\sigma = (0, 1)$ (i.e., the price is $x = 10$), the first order conditions $\partial K(\sigma,\theta)/\partial a = \partial K(\sigma,\theta)/\partial b = 0$ imply $\phi_0(10) = \phi_\theta(10) = a - b10$, and any $(a, b) \in \Theta$ on the segment AB in Figure 1 minimizes $K(\sigma, \cdot)$. These minimizers, however, lie to the right of the dashed line, where it is *not* optimal to set a price of 10. Thus, $\sigma = (0, 1)$ is not an equilibrium. A similar argument establishes that $\sigma = (1, 0)$ is not an equilibrium: If it were, the minimizer would be at D, where it is in fact not optimal to choose a price of 2.

Finally, consider mixed strategies. Because both first order conditions cannot hold simultaneously, the parameter value that minimizes $K(\sigma, \theta)$ lies on the boundary of Θ. A bit of algebra shows that, for any totally mixed σ, there is a unique minimizer $\theta_\sigma = (a_\sigma, b_\sigma)$ characterized as follows. If $\sigma_2 \leq 3/4$, the minimizer is on the segment BC:

[23] In particular, the deterministic part of the demand function can have any functional form provided it passes through $(2, \phi_o(2))$ and $(10, \phi_0(10))$.

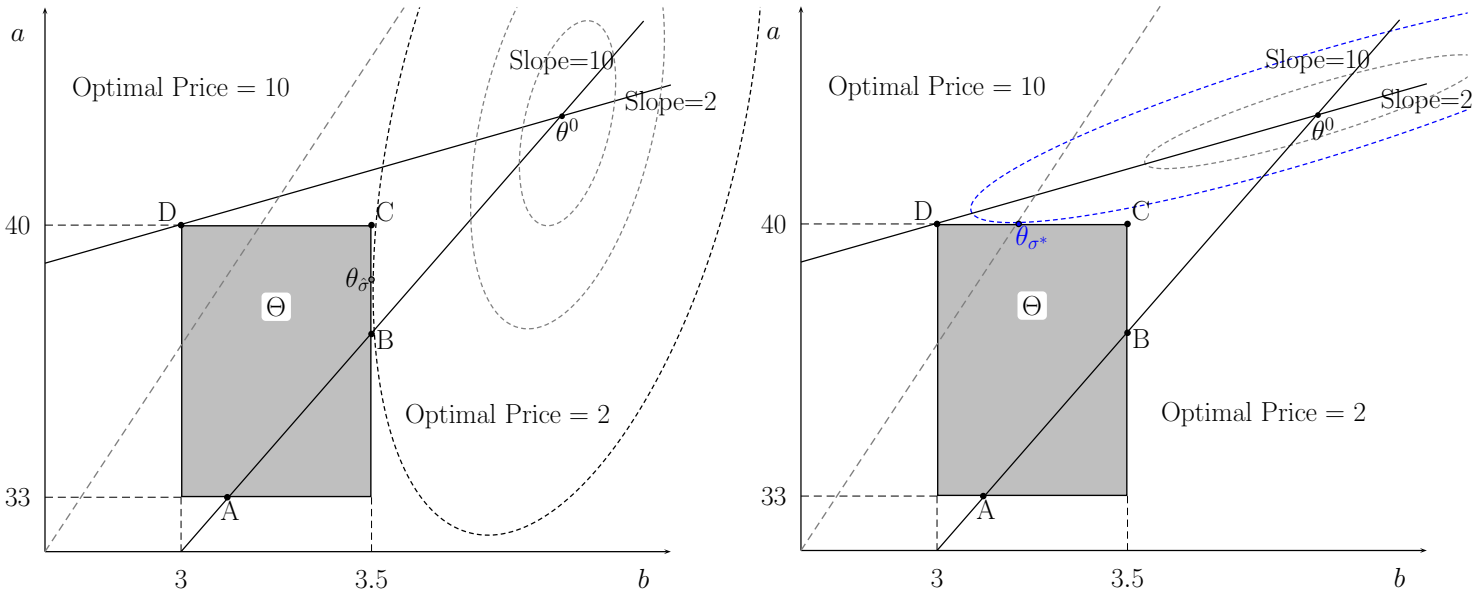

Figure 1: Monopolist with misspecified demand function.
Left panel: the parameter value that minimizes the wKLD function given strategy $\hat{\sigma}$ is $\theta_{\hat{\sigma}}$. Right panel: σ^* is a Berk-Nash equilibrium (σ^* is optimal given θ_{σ^*}—because θ_{σ^*} lies on the indifference line—and θ_{σ^*} minimizes the wKLD function given σ^*).

$b_\sigma = 3.5$ and $a_\sigma = 4\sigma_2 + 37$ solves $\partial K(\sigma, \theta)/\partial a = 0$. The left panel of Figure 1 depicts an example where the unique minimizer $\theta_{\hat{\sigma}}$ under strategy $\hat{\sigma}$ is given by the tangency between the contour lines of $K(\hat{\sigma}, \cdot)$ and the feasible set Θ.[24] If $\sigma_2 \in [3/4, 15/16]$, then $\theta_\sigma = C$ is the northeast vertex of Θ. Finally, if $\sigma_2 > 15/16$, the minimizer is on the segment DC: $a_\sigma = 40$ and $b_\sigma = (380 - 368\sigma_2)/(100 - 96\sigma_2)$ solves $\partial K(\sigma, \theta)/\partial b = 0$.

Because the monopolist mixes, optimality requires that the equilibrium belief θ_σ lies on the dashed line. The unique Berk-Nash equilibrium is $\sigma^* = (35/36, 1/36)$, and its supporting belief, $\theta_{\sigma^*} = (40, 10/3)$, is given by the intersection of the dashed line and the segment DC, as depicted in the right panel of Figure 1. It is *not* the case, however, that the equilibrium belief about the mean of Y is correct. Thus, an approach that had focused on fitting the mean, rather than minimizing K, would have led to the wrong conclusion.[25] □

[24] It can be shown that $K(\sigma, \theta) = \left(\theta - \theta^0\right)' M_\sigma \left(\theta - \theta^0\right)$, where M_σ is a weighting matrix that depends on σ. In particular, the contour lines of $K(\sigma, \cdot)$ are ellipses.

[25] The example also illustrates the importance of mixed strategies for existence of Berk-Nash equilibrium, even in single-agent settings. As an antecedent, Esponda and Pouzo (2011) argue that this is the reason why mixed strategy equilibrium cannot be purified in a voting application.

Example 2.2, continued from pg. 7. *Nonlinear* taxation. For any pure strategy x and parameter value $\theta \in \Theta^A = \mathbb{R}$ (model A) or $\theta \in \Theta^B = \mathbb{R}^2$ (model B), the wKLD function $K^j(x,\theta)$ for model $j \in \{A,B\}$ equals

$$E\Big[\ln \frac{Q(T \mid Z)p(Z-x)}{Q^j_\theta(T \mid Z)p(Z-x)} \mid X = x\Big] = \begin{cases} -\frac{1}{2}E\left[\left(\tau(Z)/Z - \theta^A\right)^2 \mid X = x\right] + C_A & \text{(model A)} \\ -\frac{1}{2}E\left[\left(\tau(Z) - \theta^B_1 - \theta^B_2 Z\right)^2 \mid X = x\right] + C_B & \text{(model B)} \end{cases}$$

where E denotes the true conditional expectation and C_A and C_B are constants.

For model A, $\theta^A(x) = E\left[\tau(x+W)/(x+W)\right]$ is the unique parameter that minimizes $K^A(x,\cdot)$.[26] Intuitively, the agent believes that the expected marginal tax rate is equal to the true expected average tax rate. For model B,

$$\theta^B_2(x) = Cov(\tau(x+W), x+W)/Var(x+W) = E\left[\tau'(x+W)\right],$$

where the second equality follows from Stein's lemma (Stein (1972)), provided that τ is differentiable. Intuitively, the agent believes that the marginal tax rate is constant and given by the true expected marginal tax rate.[27]

We now compare equilibrium under these two models with the case in which the agent has correct beliefs and chooses an optimal strategy x^{opt} that maximizes $x - E[\tau(x+W)] - c(x)$. In contrast, a strategy x^j_* is a Berk-Nash equilibrium of model j if and only if $x = x^j_*$ maximizes $x - \theta^j(x^j_*)x - c(x)$.

For example, suppose that the cost of effort and true tax schedule are both smooth functions, increasing and convex (e.g., taxes are progressive) and that $\mathbb{X} \subset \mathbb{R}$ is a compact interval. Then first order conditions are sufficient for optimality, and x^{opt} is the unique solution to $1 - E[\tau'(x^{opt}+W)] = c'(x^{opt})$. Moreover, the unique Berk-Nash equilibrium solves $1 - E\left[\tau(x^A_*+W)/(x^A_*+W)\right] = c'(x^A_*)$ for model A and $1 - E[\tau'(x^B_*+W)] = c'(x^B_*)$ for model B. In particular, effort in model B is optimal, $x^B_* = x^{opt}$. Intuitively, the agent has correct beliefs about the true expected marginal tax rate at her equilibrium choice of effort, and so she has the right incentives on the margin, despite believing incorrectly that the marginal tax rate is constant. In contrast, effort is higher than optimal in model A, $x^A_* > x^{opt}$. Intuitively, the agent

[26] We use W to denote the random variable that takes on realizations ω.

[27] By linearity and normality, the minimizers of $K^B(x,\cdot)$ coincide with the OLS estimands. We assume normality for tractability, although the framework allows for general distributional assumptions. There are other tractable distributions; for example, the minimizer of wKLD under the Laplace distribution corresponds to the estimates of a median (not a linear) regression.

believes that the expected marginal tax rate equals the true expected average tax rate, which is lower than the true expected marginal tax rate in a progressive system.□

Example 2.3, continued from pg. 7. *Regression to the mean.* Since optimal strategies are characterized by a cutoff, we let $\sigma \in \mathbb{R}$ represent the strategy where the instructor praises an initial performance if it is above σ and criticizes it otherwise. The wKLD function for any $\theta \in \Theta = \mathbb{R}^2$ is

$$K(\sigma, \theta) = \int_{-\infty}^{\sigma} E\left[\ln \frac{\varphi(S_2)}{\varphi(S_2 - (\theta_C + s_1))}\right] \varphi(s_1) ds_1 + \int_{\sigma}^{\infty} E\left[\ln \frac{\varphi(S_2)}{\varphi(S_2 - (\theta_P + s_1))}\right] \varphi(s_1) ds_1,$$

where φ is the density of $N(0,1)$ and E denotes the true expectation. For each σ, the unique parameter vector that minimizes $K(\sigma, \cdot)$ is

$$\theta_C(\sigma) = E\left[S_2 - S_1 \mid S_1 < \sigma\right] = -E\left[S_1 \mid S_1 < \sigma\right] > 0$$

and, similarly, $\theta_P(\sigma) = -E\left[S_1 \mid S_1 > \sigma\right] < 0$. Intuitively, the instructor is critical for performances below a threshold and, therefore, the mean performance conditional on a student being criticized is lower than the unconditional mean performance; thus, a student who is criticized delivers a better next performance in expectation. Similarly, a student who is praised delivers a worse next performance in expectation.

The instructor who follows a strategy cutoff σ believes, after observing initial performance $s_1 > 0$, that her expected payoff is $s_1 + \theta_C(\sigma) - \kappa s_1$ if she criticizes and $s_1 + \theta_P(\sigma)$ if she praises. By optimality, the cutoff makes her indifferent between praising and criticizing. Thus, $\sigma^* = (1/\kappa)\left(\theta_C(\sigma^*) - \theta_P(\sigma^*)\right) > 0$ is the unique equilibrium cutoff. An instructor who ignores regression to the mean has incorrect beliefs about the influence of her feedback on the student's performance: She is excessively critical in equilibrium because she incorrectly believes that criticizing a student improves performance and that praising a student worsens it. Moreover, as the reputation cost $\kappa \to 0$, meaning that instructors care only about performance and not about lying, $\sigma^* \to \infty$: instructors only criticize (as in Tversky and Kahneman's (1973) story). □

3 Relationship to other solution concepts

We show that Berk-Nash equilibrium includes several solution concepts (both standard and boundedly rational) as special cases.

3.1 Properties of games

CORRECTLY-SPECIFIED GAMES. In Bayesian statistics, a model is correctly specified if the support of the prior includes the true data generating process. The extension to single-agent decision problems is straightforward. In games, however, we must account for the fact that the objective distribution over consequences (i.e., the true model) depends on the strategy profile.[28]

Definition 2. A game is **correctly specified given $\boldsymbol{\sigma}$** if, for all $i \in I$, there exists $\theta^i \in \Theta^i$ such that $Q^i_{\theta^i}(\cdot \mid s^i, x^i) = Q^i_\sigma(\cdot \mid s^i, x^i)$ for all for all $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$; otherwise, the game is **misspecified given $\boldsymbol{\sigma}$**. A game is **correctly specified** if it is correctly specified for all σ; otherwise, it is **misspecified**.

IDENTIFICATION. From the player's perspective, what matters is identification of the distribution over consequences Q^i_θ, not the parameter θ. If the model is correctly specified, then the true Q^i_σ is trivially identified. Of course, this is not true if the model is misspecified, because the true distribution will never be learned. But we want a definition that captures the same spirit: If two distributions are judged to be equally a best fit (given the true distribution), then we want these two distributions to be identical; otherwise, we cannot identify which distribution is a best fit. The fact that players take actions introduces an additional nuance to the definition of identification. We can ask for identification of the distribution over consequences either for those actions that are taken by the player (i.e., on the path of play) or for all actions (i.e., on and off the path).

Definition 3. A game is **weakly identified given $\boldsymbol{\sigma}$** if, for all $i \in I$: if $\theta^i_1, \theta^i_2 \in \Theta^i(\sigma)$, then $Q^i_{\theta^i_1}(\cdot \mid s^i, x^i) = Q^i_{\theta^i_2}(\cdot \mid s^i, x^i)$ for all $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$ such that $\sigma^i(x^i \mid s^i) > 0$ (recall that p_{S^i} has full support). If the condition is satisfied for *all* $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$, then we say that the game is **strongly identified given $\boldsymbol{\sigma}$**. A game is [**weakly or strongly**] **identified** if it is [weakly or strongly] identified for all σ.

A correctly specified game is weakly identified. Also, two games that are identical except for their feedback may differ in terms of being correctly specified or identified.

[28] It would be more precise to say that the game is correctly specified *in steady state*.

3.2 Relationship to Nash and self-confirming equilibrium

The next result shows that Berk-Nash equilibrium is equivalent to Nash equilibrium when the game is both correctly specified and strongly identified.

Proposition 1. *(i) Suppose that the game is correctly specified given σ and that σ is a Nash equilibrium of its objective game. Then σ is a Berk-Nash equilibrium of the (objective and subjective) game; (ii) Suppose that σ is a Berk-Nash equilibrium of a game that is correctly specified and strongly identified given σ. Then σ is a Nash equilibrium of the corresponding objective game.*

Proof. (i) Let σ be a Nash equilibrium and fix any $i \in I$. Then σ^i is optimal given Q^i_σ. Because the game is correctly specified given σ, there exists $\theta^i_* \in \Theta^i$ such that $Q^i_{\theta^i_*} = Q^i_\sigma$ and, therefore, by Lemma 1, $\theta^i_* \in \Theta^i(\sigma)$. Thus, σ^i is also optimal given $Q^i_{\theta^i_*}$ and $\theta^i_* \in \Theta^i(\sigma)$, so that σ is a Berk-Nash equilibrium. (ii) Let σ be a Berk-Nash equilibrium and fix any $i \in I$. Then σ^i is optimal given $\bar{Q}^i_{\mu^i}$, for some $\mu^i \in \Delta(\Theta^i(\sigma))$. Because the game is correctly specified given σ, there exists $\theta^i_* \in \Theta^i$ such that $Q^i_{\theta^i_*} = Q^i_\sigma$ and, therefore, by Lemma 1, $\theta^i_* \in \Theta^i(\sigma)$. Moreover, because the game is strongly identified given σ, any $\hat{\theta}^i \in \Theta^i(\sigma)$ satisfies $Q^i_{\hat{\theta}^i} = Q^i_{\theta^i_*} = Q^i_\sigma$. Then σ^i is also optimal given Q^i_σ. Thus, σ is a Nash equilibrium. □

Example 2.4, continued from pg. 8. *Monetary policy*. Fix a strategy x^P_* for the public. Note that $U = u^* - \lambda(x^G - x^P_* + \varepsilon_e) + \varepsilon_U$, whereas the government believes $U = \theta_1 - \theta_2(x^G + \varepsilon_e) + \varepsilon_U$. Thus, by choosing $\theta^* \in \Theta = \mathbb{R}^2$ such that $\theta^*_1 = u^* + \lambda x^P_*$ and $\theta^*_2 = \lambda$, it follows that the distribution over $Y = (U, e)$ parameterized by θ^* coincides with the objective distribution given x^P_*. So, despite appearances, the game is correctly specified given x^P_*. Moreover, since $Var(\varepsilon_e) > 0$, θ^* is the unique minimizer of the wKLD function given x^P_*. Because there is a unique minimizer, then the game is strongly identified given x^P_*. Since these properties hold for all x^P_*, Proposition 1 implies that Berk-Nash equilibrium is equivalent to Nash equilibrium. Thus, the equilibrium policies are the same whether or not the government realizes that unemployment is driven by surprise, not actual, inflation.[29] □

[29] Sargent (1999) derived this result for a government doing OLS-based learning (a special case of our example when errors are normal). We assumed linearity for simplicity, but the result is true for the more general case with true unemployment $U = f^U(x^G, x^P, \omega)$ and subjective model $f^U_\theta(x^G, x^P, \omega)$ if, for all x^P, there exists θ such that $f^U(x^G, x^P, \omega) = f^U_\theta(x^G, x^P, \omega)$ for all (x^P, ω).

The next result shows that a Berk-Nash equilibrium is a self-confirming equilibrium (SCE) in games that are correctly specified, but not necessarily strongly identified.[30]

Proposition 2. *Suppose that the game is correctly specified given σ, and that σ is a Berk-Nash equilibrium. Then σ is also a self-confirming equilibrium.*[31]

Proof. Fix any $i \in I$ and let $\hat{\theta}^i$ be in the support of μ^i, where μ^i is player i's belief supporting the Berk-Nash equilibrium strategy σ^i. Because the game is correctly specified given σ, there exists $\theta^i_* \in \Theta^i$ such that $Q^i_{\theta^i_*} = Q^i_\sigma$ and, therefore, by Lemma 1, $K^i(\sigma, \theta^i_*) = 0$. Thus, it must also be that $K^i(\sigma, \hat{\theta}^i) = 0$. By Lemma 1, it follows that $Q^i_{\hat{\theta}^i}(\cdot \mid s^i, x^i) = Q^i_\sigma(\cdot \mid s^i, x^i)$ for all (s^i, x^i) such that $\sigma^i(x^i \mid s^i) > 0$. In particular, σ^i is optimal given $Q^i_{\hat{\theta}^i}$, and $Q^i_{\hat{\theta}^i}$ satisfies the desired self-confirming restriction. □

For games that are *not* correctly specified, beliefs can be incorrect on the equilibrium path, and so a Berk-Nash equilibrium is not necessarily Nash or SCE.

3.3 Relationship to fully cursed and ABEE

An **analogy-based game** satisfies the following four properties: (i) *States and information structure*: The state space Ω is finite with distribution $p_\Omega \in \Delta(\Omega)$. In addition, for each i, there is a partition $\mathcal{S}^i$ of Ω, and the element of $\mathcal{S}^i$ that contains ω (i.e., the signal of player i in state ω) is denoted by $s^i(\omega)$;[32] (ii) *Perfect feedback*: For each i, $f^i(x, \omega) = (x^{-i}, \omega)$ for all (x, ω); (iii) *Analogy partition*: For each i, there exists a partition of Ω, denoted by $\mathcal{A}^i$, and the element of $\mathcal{A}^i$ that contains ω is denoted by $\alpha^i(\omega)$; (iv) *Conditional independence*: $(Q^i_{\theta^i})_{\theta^i \in \Theta^i}$ is the set of all joint probability distributions over $\mathbb{X}^{-i} \times \Omega$ that satisfy

$$Q^i_{\theta^i}\left(x^{-i}, \omega \mid s^i(\omega'), x^i\right) = Q^i_{\Omega, \theta^i}(\omega \mid s^i(\omega')) Q^i_{\mathbb{X}^{-i}, \theta^i}(x^{-i} \mid \alpha^i(\omega)).$$

[30] A strategy profile σ is a SCE if, for all players $i \in I$, σ^i is optimal given $\hat{Q}^i_\sigma$, where $\hat{Q}^i_\sigma(\cdot \mid s^i, x^i) = Q^i_\sigma(\cdot \mid s^i, x^i)$ for all (s^i, x^i) such that $\sigma^i(x^i \mid s^i) > 0$. This definition is slightly more general than the typical one, e.g., Dekel et al. (2004), because it does not restrict players to believe that consequences are driven by other players' strategies.

[31] A converse does not necessarily hold for a fixed game. The reason is that the definition of SCE does not impose any restrictions on off-equilibrium beliefs, while a particular subjective game may impose ex-ante restrictions on beliefs. The following converse, however, does hold: For any σ that is an SCE, there exists a game that is correctly specified for which σ is a Berk-Nash equilibrium.

[32] This assumption is made to facilitate comparison with Jehiel and Koessler's (2008) ABEE.

In other words, every player i believes that x^{-i} and ω are independent conditional on the analogy partition. For example, if $\mathcal{A}^i = \mathcal{S}^i$ for all i, then each player believes that the actions of other players are independent of the state, conditional on their own private information.

Definition 4. (Jehiel and Koessler, 2008) A strategy profile σ is an analogy-based expectation equilibrium (ABEE) if for all $i \in I$, $\omega \in \Omega$, and x^i such that $\sigma^i(x^i \mid s^i(\omega)) > 0$, $x^i \in \arg\max_{\bar{x}^i \in \mathbb{X}^i} \sum_{\omega' \in \Omega} p_{\Omega|S^i}(\omega' \mid s^i(\omega)) \sum_{x^{-i} \in \mathbb{X}^{-i}} \bar{\sigma}^{-i}(x^{-i} \mid \omega') \pi^i(\bar{x}^i, x^{-i}, \omega')$, where $\bar{\sigma}^{-i}(x^{-i} \mid \omega') = \sum_{\omega'' \in \Omega} p_{\Omega|\mathcal{A}^i}(\omega'' \mid \alpha^i(\omega')) \prod_{j \neq i} \sigma^j(x^j \mid s^j(\omega''))$.

Proposition 3. *In an analogy-based game, σ is a Berk-Nash equilibrium if and only if it is an ABEE.*

Proof. See the Appendix. □

As mentioned by Jehiel and Koessler (2008), ABEE is equivalent to Eyster and Rabin's (2005) fully cursed equilibrium in the special case where $\mathcal{A}^i = \mathcal{S}^i$ for all i. In particular, Proposition 3 provides a misspecified-learning foundation for these two solution concepts. Jehiel and Koessler (2008) discuss an alternative foundation for ABEE, where players receive coarse feedback aggregated over past play and multiple beliefs are consistent with this feedback. Under this different feedback structure, ABEE can be viewed as a natural selection of the set of SCE.

Example 2.5, continued from pg. 8. ***Trade with adverse selection.*** In Online Appendix A, we show that x^* is a Berk-Nash equilibrium price if and only if $x = x^*$ maximizes an **equilibrium belief function** $\Pi(x, x^*)$ which represents the belief about expected profit from choosing any price x under a steady-state x^*. The equilibrium belief function depends on the feedback/misspecification assumptions, and we discuss the following four cases:

$$\begin{aligned}
\Pi^{NE}(x) &= \Pr(A \leq x)\left(E\left[V \mid A \leq x\right] - x\right) \\
\Pi^{CE}(x) &= \Pr(A \leq x)\left(E\left[V\right] - x\right) \\
\Pi^{BE}(x, x^*) &= \Pr(A \leq x)\left(E\left[V \mid A \leq x^*\right] - x\right) \\
\Pi^{ABEE}(x) &= \sum_{j=1}^{k} \Pr(V \in \mathbb{V}_j)\left\{\Pr(A \leq x \mid V \in \mathbb{V}_j)\left(E\left[V \mid V \in \mathbb{V}_j\right] - x\right)\right\}.
\end{aligned}$$

The first case, Π^{NE}, is the benchmark case in which beliefs are correct. The second case, Π^{CE}, corresponds to perfect feedback and subjective model $\Theta = \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$, as described in page 8. This is an example of an analogy-based game with single analogy class $\mathbb{V}$. The buyer learns the true marginal distributions of A and V and believes the joint distribution equals the product of the marginal distributions. Berk-Nash coincides with fully cursed equilibrium. The third case, Π^{BE}, has the same misspecified model as the second case, but assumes partial feedback, in the sense that the ask price a is always observed but the valuation v is only observed if there is trade. The equilibrium price x^* affects the sample of valuations observed by the buyer and, therefore, her beliefs. Berk-Nash coincides with naive behavioral equilibrium.

The last case, Π^{ABEE}, corresponds to perfect feedback and the following misspecification: Consider a partition of $\mathbb{V}$ into k "analogy classes" $(\mathbb{V}_j)_{j=1,...,k}$. The buyer believes that (A, V) are independent conditional on $V \in \mathbb{V}_i$, for each $i = 1, ..., k$. The parameter set is $\Theta_A = \times_j \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$, where, for a value $\theta = (\theta_1, ..., \theta_k, \theta_{\mathbb{V}}) \in \Theta_A$, $\theta_{\mathbb{V}}$ parameterizes the marginal distribution over $\mathbb{V}$ and, for each $j = 1, ..., k$, $\theta_j \in \Delta(\mathbb{A})$ parameterizes the distribution over $\mathbb{A}$ conditional on $V \in \mathbb{V}_j$. Berk-Nash coincides with the ABEE of the game with analogy classes $(\mathbb{V}_j)_{j=1,...,k}$.[33] □

4 Equilibrium foundation

We provide a learning foundation for equilibrium. We follow Fudenberg and Kreps (1993) in considering games with (slightly) perturbed payoffs because, as they highlight in the context of providing a learning foundation for mixed-strategy Nash equilibrium, behavior need not be continuous in beliefs without perturbations. Thus, even if beliefs were to converge, behavior need not settle down in the unperturbed game. Perturbations guarantee that if beliefs converge, then behavior also converges.

4.1 Perturbed game

A **perturbation structure** is a tuple $\mathcal{P} = \langle \Xi, P_\xi \rangle$, where: $\Xi = \times_{i \in I} \Xi^i$ and $\Xi^i \subseteq \mathbb{R}^{\#\mathbb{X}^i}$ is a set of payoff perturbations for each action of player i; $P_\xi = (P_{\xi^i})_{i \in I}$, where $P_{\xi^i} \in \Delta(\Xi^i)$ is a distribution over payoff perturbations of player i that is absolutely continuous with respect to the Lebesgue measure, satisfies $\int_{\Xi^i} ||\xi^i|| P_\xi(d\xi^i) < \infty$,

[33] In Online Appendix A, we also consider the case of ABEE with partial feedback.

and is independent from the perturbations of other players. A **perturbed game** $\mathcal{G}^{\mathcal{P}} = \langle \mathcal{G}, \mathcal{P} \rangle$ is composed of a game $\mathcal{G}$ and a perturbation structure $\mathcal{P}$. The timing of a perturbed game $\mathcal{G}^{\mathcal{P}}$ coincides with the timing of $\mathcal{G}$, except for two differences. First, before taking an action, each player not only observes her signal s^i but also privately observes a vector of own-payoff perturbations $\xi^i \in \Xi^i$, where $\xi^i(x^i)$ denotes the perturbation for action x^i. Second, her payoff given action x^i and consequence y^i is $\pi^i(x^i, y^i) + \xi^i(x^i)$.

A strategy σ^i for player i is **optimal in the perturbed game** given $\mu^i \in \Delta(\Theta^i)$ if, for all $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$, $\sigma^i(x^i \mid s^i) = P_\xi\left(\xi^i : x^i \in \Psi^i(\mu^i, s^i, \xi^i)\right)$, where

$$\Psi^i(\mu^i, s^i, \xi^i) \equiv \arg\max_{x^i \in \mathbb{X}^i} E_{\bar{Q}^i_{\mu^i}(\cdot|s^i,x^i)}\left[\pi^i(x^i, Y^i)\right] + \xi^i(x^i).$$

In other words, if σ^i is an optimal strategy, then $\sigma^i(x^i \mid s^i)$ is the probability that x^i is optimal when the state is s^i and the perturbation is ξ^i, taken over all possible realizations of ξ^i. The definition of Berk-Nash equilibrium of a perturbed game $\mathcal{G}^{\mathcal{P}}$ is analogous to Definition 1, with the only difference that optimality must be required with respect to the *perturbed* game.

4.2 Learning foundation

We fix a perturbed game $\mathcal{G}^{\mathcal{P}}$ and assume that players repeatedly play the corresponding objective game at each $t = 0, 1, 2, ...$, where the time-t state and signals, (ω_t, s_t), and perturbations ξ_t, are independently drawn every period from the same distribution p and P_ξ, respectively. In addition, each player i has a prior μ_0^i with full support over her (finite-dimensional) parameter set, Θ^i.[34] At the end of every period t, each player uses Bayes' rule and the information obtained in all past periods (her own signals, actions, and consequences) to update beliefs. Players believe that they face a stationary environment and myopically maximize the current period's expected payoff.

Let $\Delta^0(\Theta^i)$ denote the set of probability distributions on Θ^i with full support. Let $B^i : \Delta^0(\Theta^i) \times \mathbb{S}^i \times \mathbb{X}^i \times \mathbb{Y}^i \to \Delta^0(\Theta^i)$ denote the Bayesian operator of player i: for all

[34] We restrict attention to parametric models (i.e., finite-dimensional parameter spaces) because, otherwise, Bayesian updating need not converge to the truth for most priors and parameter values even in correctly specified statistical settings (Freedman (1963), Diaconis and Freedman (1986)).

$A \subseteq \Theta$ Borel measurable and all $(\mu^i, s^i, x^i, y^i) \in \Delta^0(\Theta^i) \times \mathbb{S}^i \times \mathbb{X}^i \times \mathbb{Y}^i$,

$$B^i(\mu^i, s^i, x^i, y^i)(A) = \frac{\int_A Q^i_{\theta^i}(y^i \mid s^i, x^i)\mu^i(d\theta)}{\int_\Theta Q^i_{\theta^i}(y^i \mid s^i, x^i)\mu^i(d\theta)}.$$

Bayesian updating is well defined by Assumption 1.[35] Because players believe they face a stationary environment with i.i.d. perturbations, it is without loss of generality to restrict player i's behavior at time t to depend on $(\mu^i_t, s^i_t, \xi^i_t)$.

Definition 5. A **policy** of player i is a sequence of functions $\phi^i = (\phi^i_t)_t$, where $\phi^i_t : \Delta(\Theta^i) \times \mathbb{S}^i \times \Xi^i \to \mathbb{X}^i$. *A* policy ϕ^i is **optimal** if $\phi^i_t \in \Psi^i$ for all t. A policy profile $\phi = (\phi^i)_{i \in I}$ *is* **optimal** if ϕ^i is optimal for all $i \in I$.

Let $\mathbb{H} \subseteq (\mathbb{S} \times \Xi \times \mathbb{X} \times \mathbb{Y})^\infty$ denote the set of histories, where any history $h = (s_0, \xi_0, x_0, y_0, ..., s_t, \xi_t, x_t, y_t ...) \in \mathbb{H}$ satisfies the feasibility restriction: for all $i \in I$, $y^i_t = f^i(x^i_t, x^{-i}_t, \omega)$ for some $\omega \in supp(p_{\Omega|S^i}(\cdot \mid s^i_t))$ for all t. Let $\mathbf{P}^{\mu_0,\phi}$ denote the probability distribution over $\mathbb{H}$ that is induced by the priors $\mu_0 = (\mu^i_0)_{i \in I}$, and the policy profiles $\phi = (\phi^i)_{i \in I}$. Let $(\mu_t)_t$ denote the sequence of beliefs $\mu_t : \mathbb{H} \to \times_{i \in I}\Delta(\Theta^i)$ such that, for all $t \geq 1$ and all $i \in I$, μ^i_t is the posterior at time t defined recursively by $\mu^i_t(h) = B^i(\mu^i_{t-1}(h), s^i_{t-1}(h), x^i_{t-1}(h), y^i_{t-1}(h))$ for all $h \in \mathbb{H}$, where $s^i_{t-1}(h)$ is player i's signal at $t-1$ given history h, and similarly for $x^i_{t-1}(h)$ and $y^i_{t-1}(h)$.

Definition 6. The **sequence of intended strategy profiles** given policy profile $\phi = (\phi^i)_{i \in I}$ is the sequence $(\sigma_t)_t$ of random variables $\sigma_t : \mathbb{H} \to \times_{i \in I}\Delta(\mathbb{X}^i)^{\mathbb{S}^i}$ such that, for all t, all $i \in I$, and all $(x^i, s^i) \in \mathbb{X}^i \times \mathbb{S}^i$,

$$\sigma^i_t(h)(x^i \mid s^i) = P_\xi\left(\xi^i : \phi^i_t(\mu^i_t(h), s^i, \xi^i) = x^i\right). \tag{7}$$

An intended strategy profile σ_t describes how each player would behave at time t for each possible signal; it is a random variable because it depends on the players' beliefs at time t, μ_t, which in turn depend on the past history.

[35] By Assumption 1(ii)-(iii), there exists $\theta \in \Theta$ and an open ball containing it, such that $Q^i_{\theta'} > 0$ for any θ' in the ball. Thus the Bayesian operator is well-defined for any $\mu^i \in \Delta^0(\Theta^i)$. Moreover, by Assumption 1(iii), such θ's are dense in Θ, so the Bayesian operator maps $\Delta^0(\Theta^i)$ into itself.

One reasonable criteria to claim that the players' behavior stabilizes is that their intended behavior stabilizes with positive probability (cf. Fudenberg and Kreps, 1993).

Definition 7. A strategy profile σ is **stable [or strongly stable]** under policy profile ϕ if the sequence of intended strategies, $(\sigma_t)_t$, converges to σ with positive probability [or with probability one], i.e.,

$$\mathbf{P}^{\mu_0,\phi}\left(\lim_{t\to\infty}\|\sigma_t(h)-\sigma\|=0\right)>0 \text{ [or } =1\text{]}.$$

Lemma 2 says that, if behavior stabilizes to a strategy profile σ, then, for each player i, beliefs become increasingly concentrated on $\Theta^i(\sigma)$. This result extends findings from the statistics of misspecified learning (Berk (1966), Bunke and Milhaud (1998)) to a setting with active learning (i.e., players learn from data that is endogenously generated by their own actions). Three new issues arise: (i) Previous results need to be extended to the case of non-i.i.d. and endogenous data; (ii) It is not obvious that steady-state beliefs can be characterized based on steady-state behavior, independently of the path of play (Assumption 1 plays an important role here; See Section 5 for an example); (iii) We allow the wKLD function to be nonfinite so that players can believe that other players follow pure strategies.[36]

Lemma 2. *Suppose that, for a policy profile ϕ, the sequence of intended strategies, $(\sigma_t)_t$, converges to σ for all histories in a set $\mathcal{H}\subseteq\mathbb{H}$ such that $\mathbf{P}^{\mu_0,\phi}(\mathcal{H})>0$. Then, for all open sets $U^i\supseteq\Theta^i(\sigma)$, $\lim_{t\to\infty}\mu_t^i(U^i)=1$, a.s.-$\mathbf{P}^{\mu_0,\phi}$ in $\mathcal{H}$.*

Proof. See the Appendix. □

The sketch of the proof of Lemma 2 is as follows (we omit the i subscript to ease the notational burden). Consider an arbitrary $\epsilon>0$ and an open set $\Theta_\epsilon(\sigma)\subseteq\Theta$ defined as the points which are within ϵ distance of $\Theta(\sigma)$. The time t posterior over the complement of $\Theta_\epsilon(\sigma)$, $\mu_t(\Theta\setminus\Theta_\epsilon(\sigma))$, can be expressed as

$$\frac{\int_{\Theta\setminus\Theta_\epsilon(\sigma)}\prod_{\tau=0}^{t-1}Q_\theta(y_\tau\mid s_\tau,x_\tau)\mu_0(d\theta)}{\int_\Theta\prod_{\tau=0}^{t-1}Q_\theta(y_\tau\mid s_\tau,x_\tau)\mu_0(d\theta)}=\frac{\int_{\Theta\setminus\Theta_\epsilon(\sigma)}e^{tK_t(\theta)}\mu_0(d\theta)}{\int_\Theta e^{tK_t(\theta)}\mu_0(d\theta)}$$

[36] For example, if player 1 believes that player 2 plays A with probability θ and B with $1-\theta$, then the wKLD function is infinity at $\theta=1$ if player 2 plays B with positive probability.

where $K_t(\theta)$ equals minus the log-likelihood ratio, $-\frac{1}{t}\sum_{\tau=0}^{t-1}\ln\frac{Q_{\sigma\tau}(y_\tau|s_\tau,x_\tau)}{Q_\theta(y_\tau|s_\tau,x_\tau)}$. This expression and straightforward algebra implies that

$$\mu_t(\Theta\setminus\Theta_\epsilon(\sigma))\leq\frac{\int_{\Theta\setminus\Theta_\epsilon(\sigma)}e^{t(K_t(\theta)+K(\sigma,\theta_0)+\delta)}\mu_0(d\theta)}{\int_{\Theta_\eta(\sigma)}e^{t(K_t(\theta)+K(\sigma,\theta_0)+\delta)}\mu_0(d\theta)}$$

for any $\delta>0$ and $\theta_0\in\Theta(\sigma)$ and $\eta>0$ taken to be "small". Roughy speaking, the integral in the numerator in the RHS is taken over points which are "ϵ-separated" from $\Theta(\sigma)$, whereas the integral in the denominator is taken over points which are "η-close" to $\Theta(\sigma)$.

Intuitively, if $K_t(\cdot)$ behaves asymptotically like $-K(\sigma,\cdot)$, there exist sufficiently small $\delta>0$ and $\eta>0$ such that $K_t(\theta)+K(\sigma,\theta_0)+\delta$ is negative for all θ which are "ϵ-separated" from $\Theta(\sigma)$, and positive for all θ which are "η-close" to $\Theta(\sigma)$. Thus, the numerator converges to zero, whereas the denominator diverges to infinity, provided that $\Theta_\eta(\sigma)$ has positive measure under the prior.

The nonstandard part of the proof consists of establishing that $\Theta_\eta(\sigma)$ has positive measure under the prior, which relies on Assumption 1, and that indeed $K_t(\cdot)$ behaves asymptotically like $-K(\sigma,\cdot)$. By virtue of Fatou's lemma, for $\theta\in\Theta_\eta(\sigma)$ it suffices to show almost sure pointwise convergence of $K_t(\theta)$ to $-K(\sigma,\theta)$; this is done in Claim B(i) in the Appendix and relies on a LLN argument for non-iid variables. On the other hand, over $\theta\in\Theta\setminus\Theta_\epsilon(\sigma)$, we need to control the asymptotic behavior of $K_t(.)$ uniformly to be able to interchange the limit and integral. In Claims B(ii) and B(iii) in the Appendix, we establish that there exists $\alpha>0$ such that asymptotically and over $\Theta\setminus\Theta_\epsilon(\sigma)$, $K_t(\cdot)<-K(\sigma,\theta_0)-\alpha$.

While Lemma 2 implies that the *support* of posteriors converges, posteriors need not converge. We can always find, however, a subsequence of posteriors that converges. By continuity of behavior in beliefs and the assumption that players are myopic, the stable strategy profile must be statically optimal. Thus, we obtain the following characterization of the set of stable strategy profiles when players follow optimal policies.

Theorem 2. *Suppose that a strategy profile σ is stable under an optimal policy profile for a perturbed game. Then σ is a Berk-Nash equilibrium of the perturbed game.*

Proof. Let ϕ denote the optimal policy function under which σ is stable. By Lemma 2, there exists $\mathcal{H}\subseteq\mathbb{H}$ with $\mathbf{P}^{\mu_0,\phi}(\mathcal{H})>0$ such that, for all $h\in\mathcal{H}$, $\lim_{t\to\infty}\sigma_t(h)=\sigma$

and $\lim_{t\to\infty} \mu_t^i(U^i) = 1$ for all $i \in I$ and all open sets $U^i \supseteq \Theta^i(\sigma)$; for the remainder of the proof, fix any $h \in \mathcal{H}$. For all $i \in I$, compactness of $\Delta(\Theta^i)$ implies the existence of a subsequence, which we denote as $(\mu^i_{t(j)})_j$, such that $\mu^i_{t(j)}$ converges (weakly) to μ^i_∞ (the limit could depend on h). We conclude by showing, for all $i \in I$:

(i) $\mu^i_\infty \in \Delta(\Theta^i(\sigma))$: Suppose not, so that there exists $\hat{\theta}^i \in supp(\mu^i_\infty)$ such that $\hat{\theta}^i \notin \Theta^i(\sigma)$. Then, since $\Theta^i(\sigma)$ is closed (by Lemma 1), there exists an open set $U^i \supset \Theta^i(\sigma)$ with closure $\bar{U}^i$ such that $\hat{\theta}^i \notin \bar{U}^i$. Then $\mu^i_\infty(\bar{U}^i) < 1$, but this contradicts the fact that $\mu^i_\infty\left(\bar{U}^i\right) \geq \limsup_{j\to\infty} \mu^i_{t(j)}\left(\bar{U}^i\right) \geq \lim_{j\to\infty} \mu^i_{t(j)}\left(U^i\right) = 1$, where the first inequality holds because $\bar{U}^i$ is closed and $\mu^i_{t(j)}$ converges (weakly) to μ^i_∞.

(ii) σ^i is optimal for the perturbed game given $\mu^i_\infty \in \Delta(\Theta^i)$:

$$\begin{aligned}\sigma^i(x^i \mid s^i) &= \lim_{j\to\infty} \sigma^i_{t(j)}(h)(x^i|s^i) = \lim_{j\to\infty} P_\xi\left(\xi^i : x^i \in \Psi^i(\mu^i_{t(j)}, s^i, \xi^i)\right) \\ &= P_\xi\left(\xi^i : x^i \in \Psi^i(\mu^i_\infty, s^i, \xi^i)\right),\end{aligned}$$

where the second equality follows because ϕ^i is optimal and Ψ^i is single-valued, a.s.-P_{ξ^i},[37] and the third equality follows from a standard continuity argument. □

4.3 A converse result

Theorem 2 provides our main justification for Berk-Nash equilibria: any strategy profile that is not an equilibrium cannot represent limiting behavior of optimizing players. Theorem 2, however, does not imply that behavior stabilizes. It is well known that convergence is not guaranteed for Nash equilibrium, which is a special case of Berk-Nash equilibrium.[38] Thus, some assumption needs to be relaxed to prove convergence for general games. Fudenberg and Kreps (1993) show that a converse for the case of Nash equilibrium can be obtained by relaxing optimality and allowing players to make vanishing optimization mistakes.

Definition 8. A policy profile ϕ is **asymptotically optimal** if there exists a positive real-valued sequence $(\varepsilon_t)_t$ with $\lim_{t\to\infty} \varepsilon_t = 0$ such that, for all $i \in I$, all $(\mu^i, s^i, \xi^i) \in$

[37] Ψ^i is single-valued a.s.-P_{ξ^i} because the set of ξ^i such that $\#\Psi^i(\mu^i, s^i, \xi^i) > 1$ is of dimension lower than $\#\mathbb{X}^i$ and, by absolute continuity of P_{ξ^i}, this set has measure zero.

[38] Jordan (1993) shows that non-convergence is robust to the choice of initial conditions; Benaim and Hirsch (1999) replicate this finding for the perturbed version of Jordan's game. In the game-theory literature, general global convergence results have only been obtained in special classes of games—e.g. zero-sum, potential, and supermodular games (Hofbauer and Sandholm, 2002).

$\Delta(\Theta^i) \times \mathbb{S}^i \times \Xi^i$, all t, and all $x^i \in \mathbb{X}^i$,

$$E_{\bar{Q}^i_{\mu^i}(\cdot|s^i,x^i_t)}\left[\pi^i(x^i_t, Y^i)\right] + \xi^i(x^i_t) \geq E_{\bar{Q}^i_{\mu^i}(\cdot|s^i,x^i)}\left[\pi^i(x^i, Y^i)\right] + \xi^i(x^i) - \varepsilon_t$$

where $x^i_t = \phi^i_t(\mu^i, s^i, \xi^i)$.

Fudenberg and Kreps' (1993) insight is to suppose that players are convinced early on that the equilibrium strategy is the right one to play, and continue to play this strategy unless they have strong enough evidence to think otherwise. And, as they continue to play the equilibrium strategy, the evidence increasingly convinces them that it is the right thing to do. This idea, however, need not work for Berk-Nash equilibrium because beliefs may not converge if the model is misspecified (see Berk (1966) for an example). If the game is weakly identified, however, Lemma 2 and Fudenberg and Kreps' (1993) insight can be combined to obtain the following converse of Theorem 2.

Theorem 3. *Suppose that σ is a Berk-Nash equilibrium of a perturbed game that is weakly identified given σ. Then there exists a profile of priors with full support and an asymptotically optimal policy profile ϕ such that σ is strongly stable under ϕ.*

Proof. See Online Appendix B.[39] □

5 Discussion

IMPORTANCE OF ASSUMPTION 1. The following example illustrates that equilibrium may not exist and Lemma 2 fails if Assumption 1 does not hold. A single agent chooses action $x \in \{A, B\}$ and obtains outcome $y \in \{0, 1\}$. The agent's model is parameterized by $\theta = (\theta_A, \theta_B)$, where $Q_\theta(y = 1 \mid A) = \theta_A$ and $Q_\theta(y = 1 \mid B) = \theta_B$. The true model is $\theta^0 = (1/4, 3/4)$. The agent, however, is misspecified and considers only $\theta_1 = (0, 3/4)$ and $\theta_2 = (1/4, 1/4)$ to be possible, i.e., $\Theta = \{\theta_1, \theta_2\}$. In particular, Assumption 1(iii) fails for parameter value θ_1.[40] Suppose that A is uniquely optimal for parameter value θ_1 and B is uniquely optimal for θ_2 (further details about payoffs are not needed).

[39] The requirement that the priors have full support makes the statement non trivial.

[40] Assumption 1(iii) would hold if, for some $\bar{\varepsilon} > 0$, $\theta = (\varepsilon, 3/4)$ were also in Θ for all $0 < \varepsilon \leq \bar{\varepsilon}$.

A Berk-Nash equilibrium does not exist: If A is played with positive probability, then the wKLD is infinity at θ_1 (i.e., θ_1 cannot rationalize $y = 1$ given A) and θ_2 is the best fit; but then A is not optimal. If B is played with probability 1, then θ_1 is the best fit; but then B is not optimal. In addition, Lemma 2 fails: Suppose that the path of play converges to pure strategy B. The best fit given B is θ_1, but the posterior need not converge weakly to a degenerate probability distribution on θ_1; it is possible that, along the path of play, the agent tried action A and observed $y = 1$, in which case the posterior would immediately assign probability 1 to θ_2.

FORWARD-LOOKING AGENTS. In the dynamic model, we assumed that players are myopic. In Online Appendix C, we extend Theorem 2 to the case of non-myopic players who solve a dynamic optimization problem with beliefs as a state variable. A key fact used in the proof of Theorem 2 is that myopically optimal behavior is continuous in beliefs. Non-myopic optimal behavior is also continuous in beliefs, but the issue is that it may not coincide with myopic behavior in the steady state if players still have incentives to experiment. We prove the extension by requiring that the game is weakly identified, which guarantees that players have no incentives to experiment in steady state.

LARGE POPULATION MODELS. The framework assumes that there is a fixed number of players but, by focusing on stationary subjective models, rules out aspects of "repeated games" where players attempt to influence each others' play. In Online Appendix D, we adapt the equilibrium concept to settings in which there is a population of a large number of agents in the role of each player, so that agents have negligible incentives to influence each other's play.

EXTENSIVE-FORM GAMES. Our results hold for an alternative timing where player i commits to a signal-contingent plan of action (i.e., a strategy) and observes both the realized signal s^i and the consequence y^i ex post. In particular, Berk-Nash equilibrium is applicable to extensive-form games provided that players compete by choosing contingent plan of actions and know the extensive form. But the right approach is less clear if players have a misspecified view of the extensive form (for example, they may not even know the set of strategies available to them) or if players play the game sequentially (for example, we would need to define and update beliefs at each information set). The extension to extensive-form games is left for future work.[41]

[41] Jehiel (1995) considers the class of repeated alternating-move games and assumes that players only forecast a limited number of time periods into the future; see Jehiel (1998) for a learning foundation. Jehiel and Samet (2007) consider the general class of extensive form games with perfect

RELATIONSHIP TO BOUNDED RATIONALITY LITERATURE. By providing a language that makes the underlying misspecification explicit, we offer some guidance for choosing between different models of bounded rationality. For example, we could model the observed behavior of an instructor in Example 2.3 by directly assuming that she believes criticism improves performance and praise worsens it.[42] But extrapolating this observed belief to other contexts may lead to erroneous conclusions. Instead, we postulate what we think is a plausible misspecification (i.e., failure to account for regression to the mean) and then derive beliefs endogenously, as a function of the context.

We mentioned in the paper several instances of bounded rationality that can be formalized via misspecified, endogenous learning. Other examples in the literature can also be viewed as restricting beliefs using the wKLD measure, but fall outside the scope of our paper either because interactions are mediated by a price or because the problem is dynamic (we focus on the repetition of a static problem). For example, Blume and Easley (1982) and Rabin and Vayanos (2010) explicitly characterize beliefs using the limit of a likelihood function, while Bray (1982), Radner (1982), Sargent (1993), and Evans and Honkapohja (2001) focus specifically on OLS learning with misspecified models. Piccione and Rubinstein (2003), Eyster and Piccione (2013), and Spiegler (2013) study pattern recognition in dynamic settings and impose consistency requirements on beliefs that could be interpreted as minimizing the wKLD measure. In the sampling equilibrium of Osborne and Rubinstein (1998) and Spiegler (2006), beliefs may be incorrect due to learning from a limited sample, rather than from misspecified learning. Other instances of bounded rationality that do not seem naturally fitted to misspecified learning, include biases in information processing due to computational complexity (e.g., Rubinstein (1986), Salant (2011)), bounded memory (e.g., Wilson, 2003), self-deception (e.g., Bénabou and Tirole (2002), Compte and Postlewaite (2004)), or sparsity-based optimization (Gabaix (2014)).

information and assume that players simplify the game by partitioning the nodes into similarity classes. In both cases, players are required to have correct beliefs, given their limited or simplified view of the game.

[42] This assumption corresponds to a singleton set Θ, thus fixing beliefs at the outset and leaving no space for learning. This approach is common in past work that assumes that agents have a misspecified model but there is no learning about parameter values, e.g., Barberis et al. (1998).

References

Al-Najjar, N., "Decision Makers as Statisticians: Diversity, Ambiguity and Learning," *Econometrica*, 2009, *77* (5), 1371–1401.

_ **and M. Pai**, "Coarse decision making and overfitting," *Journal of Economic Theory, forthcoming*, 2013.

Aliprantis, C.D. and K.C. Border, *Infinite dimensional analysis: a hitchhiker's guide*, Springer Verlag, 2006.

Aragones, E., I. Gilboa, A. Postlewaite, and D. Schmeidler, "Fact-Free Learning," *American Economic Review*, 2005, *95* (5), 1355–1368.

Arrow, K. and J. Green, "Notes on Expectations Equilibria in Bayesian Settings," *Institute for Mathematical Studies in the Social Sciences Working Paper No. 33*, 1973.

Barberis, N., A. Shleifer, and R. Vishny, "A model of investor sentiment," *Journal of financial economics*, 1998, *49* (3), 307–343.

Battigalli, P., *Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali*, Universita Bocconi, Milano, 1987.

_ **, S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci**, "Selfconfirming equilibrium and model uncertainty," Technical Report 2012.

Bénabou, Roland and Jean Tirole, "Self-confidence and personal motivation," *The Quarterly Journal of Economics*, 2002, *117* (3), 871–915.

Benaim, M. and M.W. Hirsch, "Mixed equilibria and dynamical systems arising from fictitious play in perturbed games," *Games and Economic Behavior*, 1999, *29* (1-2), 36–72.

Benaim, Michel, "Dynamics of stochastic approximation algorithms," in "Seminaire de Probabilites XXXIII," Vol. 1709 of *Lecture Notes in Mathematics*, Springer Berlin Heidelberg, 1999, pp. 1–68.

Berk, R.H., "Limiting behavior of posterior distributions when the model is incorrect," *The Annals of Mathematical Statistics*, 1966, *37* (1), 51–58.

Bickel, Peter J, Chris AJ Klaassen, Peter J Bickel, Y Ritov, J Klaassen, Jon A Wellner, and YA'Acov Ritov, *Efficient and adaptive estimation for semiparametric models*, Johns Hopkins University Press Baltimore, 1993.

Billingsley, P., *Probability and Measure*, Wiley, 1995.

Blume, L.E. and D. Easley, "Learning to be Rational," *Journal of Economic Theory*, 1982, *26* (2), 340–351.

Bray, M., "Learning, estimation, and the stability of rational expectations," *Journal of economic theory*, 1982, *26* (2), 318–339.

Bunke, O. and X. Milhaud, "Asymptotic behavior of Bayes estimates under possibly incorrect models," *The Annals of Statistics*, 1998, *26* (2), 617–644.

Compte, Olivier and Andrew Postlewaite, "Confidence-enhanced performance," *American Economic Review*, 2004, pp. 1536–1557.

Dekel, E., D. Fudenberg, and D.K. Levine, "Learning to play Bayesian games," *Games and Economic Behavior*, 2004, *46* (2), 282–303.

Diaconis, P. and D. Freedman, "On the consistency of Bayes estimates," *The Annals of Statistics*, 1986, pp. 1–26.

Doraszelski, Ulrich and Juan F Escobar, "A theory of regular Markov perfect equilibria in dynamic stochastic games: Genericity, stability, and purification," *Theoretical Economics*, 2010, *5* (3), 369–402.

Durrett, R., *Probability: Theory and Examples*, Cambridge University Press, 2010.

Easley, D. and N.M. Kiefer, "Controlling a stochastic process with unknown parameters," *Econometrica*, 1988, pp. 1045–1064.

Esponda, I., "Behavioral equilibrium in economies with adverse selection," *The American Economic Review*, 2008, *98* (4), 1269–1291.

_ **and D. Pouzo**, "Learning Foundation for Equilibrium in Voting Environments with Private Information," *working paper*, 2011.

Esponda, I. and D. Pouzo, "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," *ArXiv 1411.1152*, November 2014.

Evans, G. W. and S. Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton University Press, 2001.

Eyster, E. and M. Rabin, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

Eyster, Erik and Michele Piccione, "An approach to asset-pricing under incomplete and diverse perceptions," *Econometrica*, 2013, *81* (4), 1483–1506.

Freedman, D.A., "On the asymptotic behavior of Bayes' estimates in the discrete case," *The Annals of Mathematical Statistics*, 1963, *34* (4), 1386–1403.

Fudenberg, D. and D. Kreps, "Learning Mixed Equilibria," *Games and Economic Behavior*, 1993, *5*, 320–367.

_ **and D.K. Levine**, "Self-confirming equilibrium," *Econometrica*, 1993, pp. 523–545.

_ **and** _ , "Steady state learning and Nash equilibrium," *Econometrica*, 1993, pp. 547–573.

_ **and** _ , *The theory of learning in games*, Vol. 2, The MIT press, 1998.

_ **and** _ , "Learning and Equilibrium," *Annual Review of Economics*, 2009, *1*, 385–420.

_ **and D.M. Kreps**, "A Theory of Learning, Experimentation, and Equilibrium in Games," Technical Report, mimeo 1988.

_ **and** _ , "Learning in extensive-form games I. Self-confirming equilibria," *Games and Economic Behavior*, 1995, *8* (1), 20–55.

Gabaix, Xavier, "A sparsity-based model of bounded rationality," *The Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

Harsanyi, J.C., "Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points," *International Journal of Game Theory*, 1973, *2* (1), 1–23.

Hirsch, M. W., S. Smale, and R. L. Devaney, *Differential Equations, Dynamical Systems and An Introduction to Chaos*, Elsevier Academic Press, 2004.

Hofbauer, J. and W.H. Sandholm, "On the global convergence of stochastic fictitious play," *Econometrica*, 2002, *70* (6), 2265–2294.

Jehiel, P., "Analogy-based expectation equilibrium," *Journal of Economic theory*, 2005, *123* (2), 81–104.

_ **and D. Samet**, "Valuation equilibrium," *Theoretical Economics*, 2007, *2* (2), 163–185.

_ **and F. Koessler**, "Revisiting games of incomplete information with analogy-based expectations," *Games and Economic Behavior*, 2008, *62* (2), 533–557.

Jehiel, Philippe, "Learning to play limited forecast equilibria," *Games and Economic Behavior*, 1998, *22* (2), 274–298.

Jehiel, Phillippe, "Limited horizon forecast in repeated alternate games," *Journal of Economic Theory*, 1995, *67* (2), 497–519.

Jordan, J. S., "Three problems in learning mixed-strategy Nash equilibria," *Games and Economic Behavior*, 1993, *5* (3), 368–386.

Kagel, J.H. and D. Levin, "The winner's curse and public information in common value auctions," *The American Economic Review*, 1986, pp. 894–920.

Kalai, E. and E. Lehrer, "Rational learning leads to Nash equilibrium," *Econometrica*, 1993, pp. 1019–1045.

Kirman, A. P., "Learning by firms about demand conditions," in R. H. Day and T. Groves, eds., *Adaptive economic models*, Academic Press 1975, pp. 137–156.

Kreps, D. M. and R. Wilson, "Sequential equilibria," *Econometrica*, 1982, pp. 863–894.

Kullback, S. and R. A. Leibler, "On Information and Sufficiency," *Annals of Mathematical Statistics*, 1951, *22* (1), 79–86.

Kushner, H. J. and G. G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, Springer Verlag, 2003.

McLennan, A., "Price dispersion and incomplete learning in the long run," *Journal of Economic Dynamics and Control*, 1984, *7* (3), 331–347.

Nyarko, Y., "Learning in mis-specified models and the possibility of cycles," *Journal of Economic Theory*, 1991, *55* (2), 416–427.

_ , "On the convexity of the value function in Bayesian optimal control problems," *Economic Theory*, 1994, *4* (2), 303–309.

Osborne, M.J. and A. Rubinstein, "Games with procedurally rational players," *American Economic Review*, 1998, *88*, 834–849.

Piccione, M. and A. Rubinstein, "Modeling the economic interaction of agents with diverse abilities to recognize equilibrium patterns," *Journal of the European economic association*, 2003, *1* (1), 212–223.

Pollard, D., *A User's Guide to Measure Theoretic Probability*, Cambridge University Press, 2001.

Rabin, M., "Inference by Believers in the Law of Small Numbers," *Quarterly Journal of Economics*, 2002, *117* (3), 775–816.

_ **and D. Vayanos**, "The gambler's and hot-hand fallacies: Theory and applications," *The Review of Economic Studies*, 2010, *77* (2), 730–778.

Radner, R., *Equilibrium Under Uncertainty*, Vol. II of *Handbook of Mathematical Economies*, North-Holland Publishing Company, 1982.

Rothschild, M., "A two-armed bandit theory of market pricing," *Journal of Economic Theory*, 1974, *9* (2), 185–202.

Rubinstein, A. and A. Wolinsky, "Rationalizable conjectural equilibrium: between Nash and rationalizability," *Games and Economic Behavior*, 1994, *6* (2), 299–311.

Rubinstein, Ariel, "Finite automata play the repeated prisoner's dilemma," *Journal of economic theory*, 1986, *39* (1), 83–96.

Salant, Y., "Procedural analysis of choice rules with applications to bounded rationality," *The American Economic Review*, 2011, *101* (2), 724–748.

Sargent, T. J., *Bounded rationality in macroeconomics*, Oxford University Press, 1993.

_ , *The Conquest of American Inflation*, Princeton University Press, 1999.

Schwartzstein, J., "Selective Attention and Learning," *working paper*, 2009.

Selten, R., "Reexamination of the perfectness concept for equilibrium points in extensive games," *International journal of game theory*, 1975, *4* (1), 25–55.

Sobel, J., "Non-linear prices and price-taking behavior," *Journal of Economic Behavior & Organization*, 1984, *5* (3), 387–396.

Spiegler, R., "The Market for Quacks," *Review of Economic Studies*, 2006, *73*, 1113–1131.

_ , "Placebo reforms," *The American Economic Review*, 2013, *103* (4), 1490–1506.

_ , "Bayesian Networks and Boundedly Rational Expectations," *Working Paper*, 2014.

Stein, Charles, "A bound for the error in the normal approximation to the distribution of a sum of dependent random variables," in "Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory" University of California Press Berkeley, Calif. 1972, pp. 583–602.

Tversky, T. and D. Kahneman, "Availability: A heuristic for judging frequency and probability," *Cognitive Psychology*, 1973, *5*, 207–232.

White, Halbert, "Maximum likelihood estimation of misspecified models," *Econometrica: Journal of the Econometric Society*, 1982, pp. 1–25.

Wilson, A., "Bounded Memory and Biases in Information Processing," *Working Paper*, 2003.

Appendix

Let $\mathbb{Z}^i = \big\{(s^i, x^i, y^i) \in \mathbb{S}^i \times \mathbb{X}^i \times \mathbb{Y}^i : y^i = f^i(x^i, x^{-i}, \omega), x^{-i} \in \mathbb{X}^{-i}, \omega \in supp(p_{\Omega|S^i}(\cdot \mid s^i))\big\}$. For all $z^i = (s^i, x^i, y^i) \in \mathbb{Z}^i$, define $\bar{P}^i_\sigma(z^i) = Q^i_\sigma(y^i \mid s^i, x^i)\sigma^i(x^i \mid s^i)p_{S^i}(s^i)$. We sometimes abuse notation and write $Q^i_\sigma(z^i) \equiv Q^i_\sigma(y^i \mid s^i, x^i)$, and similarly for $Q^i_{\theta^i}$. The following claim is used in the proofs below.

Claim A. *For all $i \in I$: **(i)** There exists $\theta^i_* \in \Theta^i$ and $\bar{K} < \infty$ such that, $\forall \sigma \in \Sigma$, $K^i(\sigma, \theta^i_*) \leq \bar{K}$; **(ii)** Fix any $\theta^i \in \Theta^i$ and $(\sigma_n)_n$ such that $Q^i_{\theta^i}(z^i) > 0$ $\forall z^i \in \mathbb{Z}^i$ and $\lim_{n\to\infty} \sigma_n = \sigma$. Then $\lim_{n\to\infty} K^i(\sigma_n, \theta^i) = K^i(\sigma, \theta^i)$; **(iii)** K^i is (jointly) lower semicontinuous: Fix any $(\theta^i_n)_n$ and $(\sigma_n)_n$ such that $\lim_{n\to\infty} \theta^i_n = \theta^i$, $\lim_{n\to\infty} \sigma_n = \sigma$. Then $\liminf_{n\to\infty} K^i(\sigma_n, \theta^i_n) \geq K(\sigma, \theta^i)$; **(iv)** Let ξ^i be a random vector in $\mathbb{R}^{\#\mathbb{X}^i}$ with absolutely continuous probability distribution P_ξ. Then, $\forall (s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$, $\mu^i \mapsto \sigma^i(\mu^i)(x^i \mid s^i) = P_\xi\big(\xi^i : x^i \in \arg\max_{\bar{x}^i \in \mathbb{X}^i} E_{\bar{Q}^i_{\mu^i}(\cdot|s^i,\bar{x}^i)}[\pi^i(\bar{x}^i, Y^i)] + \xi^i(\bar{x}^i)\big)$ is continuous.*

Proof. (i) By Assumption 1 and finiteness of $\mathbb{Z}^i$, there exist $\theta^i_* \in \Theta$ and $\alpha \in (0,1)$ such that $Q^i_{\theta^i_*}(z^i) \geq \alpha$ $\forall z^i \in \mathbb{Z}^i$. Thus, $\forall \sigma \in \Sigma$, $K(\sigma, \theta_*) \leq -E_{\bar{P}^i_\sigma}[\ln Q^i_{\theta_*}(Z^i)] \leq -\ln \alpha$.

(ii) $K^i(\sigma_n, \theta^i) - K(\sigma, \theta^i) = \sum_{z^i \in \mathbb{Z}^i} \big\{\big(\bar{P}^i_{\sigma_n}(z^i) \ln Q^i_{\sigma_n}(z^i) - \bar{P}^i_\sigma(z^i) \ln Q^i_\sigma(z^i)\big) + \big(\bar{P}^i_\sigma(z^i) - \bar{P}^i_{\sigma_n}(z^i)\big) \ln Q^i_{\theta^i}(z^i)\big\}$. The first term in the RHS converges to 0 because $\lim_{n\to\infty} \sigma_n = \sigma$, Q_σ is continuous, and $x \ln x$ is continuous $\forall x \in [0,1]$. The second term converges to 0 because $\lim_{n\to\infty} \sigma_n = \sigma$, $\bar{P}^i_\sigma$ is continuous, and $\ln Q^i_{\theta^i}(z^i) \in (-\infty, 0]$ $\forall z^i \in \mathbb{Z}^i$.

(iii) $K^i(\sigma_n, \theta^i_n) - K(\sigma, \theta^i) = \sum_{z^i \in \mathbb{Z}^i} \big\{\big(\bar{P}^i_{\sigma_n}(z^i) \ln Q^i_{\sigma_n}(z^i) - \bar{P}^i_\sigma(z^i) \ln Q^i_\sigma(z^i)\big) + \big(\bar{P}^i_\sigma(z^i) \ln Q^i_{\theta^i}(z^i) - \bar{P}^i_{\sigma_n}(z^i) \ln Q^i_{\theta^i_n}(z^i)\big)\big\}$. The first term in the RHS converges to 0 (same argument as in part (ii)). The proof concludes by showing that, $\forall z^i \in \mathbb{Z}^i$,

$$\lim \inf_{n\to\infty} -\bar{P}^i_{\sigma_n}(z^i) \ln Q^i_{\theta^i_n}(z^i) \geq -\bar{P}^i_\sigma(z^i) \ln Q^i_{\theta^i}(z^i). \tag{8}$$

Suppose $\liminf_{n\to\infty} -\bar{P}^i_{\sigma_n}(z^i) \ln Q^i_{\theta^i_n}(z^i) \leq M < \infty$ (if not, (8) holds trivially). Then either (i) $\bar{P}^i_{\sigma_n}(z^i) \to \bar{P}^i_\sigma(z^i) > 0$, in which case (8) holds with equality (by continuity of $\theta^i \mapsto Q^i_{\theta^i}$), or (ii) $\bar{P}^i_{\sigma_n}(z^i) \to \bar{P}^i_\sigma(z^i) = 0$, in which case (8) holds because its RHS is 0 (by convention that $0 \ln 0 = 0$) and its LHS is always nonnegative.

(iv) The proof is standard and, therefore, omitted. □

Proof of Lemma 1. Part (i). Note that

$$K^i(\sigma, \theta^i) \geq -\sum_{(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i} \ln\big(E_{Q^i_\sigma(\cdot|s^i,x^i)}\big[\frac{Q^i_{\theta^i}(Y^i \mid s^i, x^i)}{Q^i_\sigma(Y^i \mid s^i, x^i)}\big]\big)\sigma^i(x^i \mid s^i)p_{S^i}(s^i) = 0,$$

where the inequality follows from Jensen's inequality and the strict concavity of $\ln(\cdot)$, and it holds with equality if and only if $Q^i_{\theta^i}(\cdot \mid s^i, x^i) = Q^i_{\theta^i}(\cdot \mid s^i, x^i)\ \forall (s^i, x^i)$ such that $\sigma^i(x^i \mid s^i) > 0$ (recall that, by assumption, $p_{S^i}(s^i) > 0$).

Part (ii). $\Theta^i(\sigma)$ *is nonempty*: By Claim A(i), $\exists K < \infty$ such that the minimizers are in the constraint set $\{\theta^i \in \Theta^i : K^i(\sigma, \theta^i) \leq K\}$. Because $K^i(\sigma, \cdot)$ is continuous over a compact set, a minimum exists.

$\Theta^i(\sigma)$ *is upper hemicontinuous (uhc)*: Fix any $(\sigma_n)_n$ and $(\theta^i_n)_n$ such that $\lim_{n\to\infty} \sigma_n = \sigma$, $\lim_{n\to\infty} \theta^i_n = \theta^i$, and $\theta^i_n \in \Theta^i(\sigma_n)\ \forall n$. We show that $\theta^i \in \Theta^i(\sigma)$ (so that $\Theta^i(\cdot)$ has a closed graph and, by compactness of Θ^i, is uhc). Suppose, to obtain a contradiction, that $\theta^i \notin \Theta^i(\sigma)$. By Claim A(i), there exist $\hat{\theta}^i \in \Theta^i$ and $\varepsilon > 0$ such that $K^i(\sigma, \hat{\theta}^i) \leq K^i(\sigma, \theta^i) - 3\varepsilon$ and $K^i(\sigma, \hat{\theta}^i) < \infty$. By Assumption 1, $\exists (\hat{\theta}^i_j)_j$ with $\lim_{j\to\infty} \hat{\theta}^i_j = \hat{\theta}^i$ and, $\forall j$, $Q^i_{\hat{\theta}^i_j}(z^i) > 0\ \forall z^i \in \mathbb{Z}^i$. We show that there is an element of the sequence, $\hat{\theta}^i_J$, that "does better" than θ^i_n given σ_n, which is a contradiction. Because $K^i(\sigma, \hat{\theta}^i) < \infty$, continuity of $K^i(\sigma, \cdot)$ implies that there exists J large enough such that $\left|K^i(\sigma, \hat{\theta}^i_J) - K^i(\sigma, \hat{\theta}^i)\right| \leq \varepsilon/2$. Moreover, Claim A(ii) applied to $\theta^i = \hat{\theta}^i_J$ implies that there exists $N_{\varepsilon,J}$ such that, $\forall n \geq N_{\varepsilon,J}$, $\left|K^i(\sigma_n, \hat{\theta}^i_J) - K^i(\sigma, \hat{\theta}^i_J)\right| \leq \varepsilon/2$. Thus, $\forall n \geq N_{\varepsilon,J}$, $\left|K^i(\sigma_n, \hat{\theta}^i_J) - K^i(\sigma, \hat{\theta}^i)\right| \leq \left|K^i(\sigma_n, \hat{\theta}^i_J) - K^i(\sigma, \hat{\theta}^i_J)\right| + \left|K^i(\sigma, \hat{\theta}^i_J) - K^i(\sigma, \hat{\theta}^i)\right| \leq \varepsilon$. Therefore,

$$K^i(\sigma_n, \hat{\theta}^i_J) \leq K^i(\sigma, \hat{\theta}^i) + \varepsilon \leq K^i(\sigma, \theta^i) - 2\varepsilon. \tag{9}$$

Suppose $K^i(\sigma, \theta^i) < \infty$. By Claim A(iii), $\exists n_\varepsilon \geq N_{\varepsilon,J}$ such that $K^i(\sigma_{n_\varepsilon}, \theta^i_{n_\varepsilon}) \geq K^i(\sigma, \theta^i) - \varepsilon$. This result and expression (9), imply $K^i(\sigma_{n_\varepsilon}, \hat{\theta}^i_J) \leq K(\sigma_{n_\varepsilon}, \theta^i_{n_\varepsilon}) - \varepsilon$. But this contradicts $\theta^i_{n_\varepsilon} \in \Theta^i(\sigma_{n_\varepsilon})$. Finally, if $K^i(\sigma, \theta^i) = \infty$, Claim A(iii) implies that $\exists n_\varepsilon \geq N_{\varepsilon,J}$ such that $K^i(\sigma_{n_\varepsilon}, \theta^i_{n_\varepsilon}) \geq 2K$, where K is the bound defined in Claim A(i). But this also contradicts $\theta^i_{n_\varepsilon} \in \Theta^i(\sigma_{n_\varepsilon})$.

$\Theta^i(\sigma)$ *is compact*: As shown above, $\Theta^i(\cdot)$ has a closed graph, and so $\Theta^i(\sigma)$ is a closed set. Compactness of $\Theta^i(\sigma)$ follows from compactness of Θ^i. □

Proof of Theorem 1. We prove the result in two parts. *Part 1.* We show existence of equilibrium in the perturbed game (defined in Section 4.1). Let $\Gamma : \times_{i\in I}\Delta(\Theta^i) \rightrightarrows \times_{i\in I}\Delta(\Theta^i)$ be a correspondence such that, $\forall \mu = (\mu^i)_{i\in I} \in \times_{i\in I}\Delta(\Theta^i)$, $\Gamma(\mu) = \times_{i\in I}\Delta\left(\Theta^i(\sigma(\mu))\right)$, where $\sigma(\mu) = (\sigma^i(\mu^i))_{i\in I} \in \Sigma$ and is defined as

$$\sigma^i(\mu^i)(x^i|s^i) = P_\xi\left(\xi^i : x^i \in \arg\max_{\bar{x}^i \in \mathbb{X}^i} E_{\bar{Q}^i_{\mu^i}(\cdot|s^i,\bar{x}^i)}\left[\pi^i(\bar{x}^i, Y^i)\right] + \xi^i(\bar{x}^i)\right) \tag{10}$$

$\forall (x^i, s^i) \in \mathbb{X}^i \times \mathbb{S}^i$. Note that if $\exists \mu_* \in \times_{i\in I}\Delta(\Theta^i)$ such that $\mu_* \in \Gamma(\mu_*)$, then $\sigma_* \equiv (\sigma^i(\mu_*^i))_{i\in I}$ is an equilibrium of the perturbed game. We show that such μ_* exists by checking the conditions of the Kakutani-Fan-Glicksberg fixed-point theorem: (i) $\times_{i\in I}\Delta(\Theta^i)$ is compact, convex and locally convex Hausdorff: The set $\Delta(\Theta^i)$ is convex, and since Θ^i is compact $\Delta(\Theta^i)$ is also compact under the weak topology (Aliprantis and Border (2006), Theorem 15.11). By Tychonoff's theorem, $\times_{i\in I}\Delta(\Theta^i)$ is compact too. Finally, the set is also locally convex under the weak topology;[43] (ii) Γ has convex, nonempty images: It is clear that $\Delta\left(\Theta^i(\sigma(\mu))\right)$ is convex valued $\forall \mu$. Also, by Lemma 1, $\Theta^i(\sigma(\mu))$ is nonempty $\forall \mu$; (iii) Γ has a closed graph: Let $(\mu_n, \hat{\mu}_n)_n$ be such that $\hat{\mu}_n \in \Gamma(\mu_n)$ and $\mu_n \to \mu$ and $\hat{\mu}_n \to \hat{\mu}$ (under the weak topology). By Claim A(iv), $\mu^i \mapsto \sigma^i(\mu^i)$ is continuous. Thus, $\sigma_n \equiv (\sigma^i(\mu_n^i))_{i\in I} \to \sigma \equiv (\sigma^i(\mu^i))_{i\in I}$. By Lemma 1, $\sigma \mapsto \Theta^i(\sigma)$ is uhc; thus, by Theorem 17.13 in Aliprantis and Border (2006), $\sigma \mapsto \times_{i\in I}\Delta\left(\Theta^i(\sigma)\right)$ is also uhc. Therefore, $\hat{\mu} \in \times_{i\in I}\Delta\left(\Theta^i(\sigma)\right) = \Gamma(\mu)$.

Part 2. Fix a sequence of perturbed games indexed by the probability of perturbations $(P_{\xi,n})_n$. By Part 1, there is a corresponding sequence of fixed points $(\mu_n)_n$, such that $\mu_n \in \times_{i\in I}\Delta\left(\Theta^i(\sigma_n)\right)\ \forall n$, where $\sigma_n \equiv (\sigma^i(\mu_n^i, P_{\xi,n}))_{i\in I}$ (see equation (10), where we now explicitly account for dependance on $P_{\xi,n}$). By compactness, there exist subsequences of $(\mu_n)_n$ and $(\sigma_n)_n$ that converge to μ and σ, respectively. Since $\sigma \mapsto \times_{i\in I}\Delta\left(\Theta^i(\sigma)\right)$ is uhc, then $\mu \in \times_{i\in I}\Delta\left(\Theta^i(\sigma)\right)$. We now show that if we choose $(P_{\xi,n})_n$ such that, $\forall \varepsilon > 0$, $\lim_{n\to\infty} P_{\xi,n}\left(\|\xi_n^i\| \geq \varepsilon\right) = 0$, then σ is optimal given μ in the unperturbed game—this establishes existence of equilibrium in the unperturbed game. Suppose not, so that there exist $i, s^i, x^i, \hat{x}^i$, and $\varepsilon > 0$ such that $\sigma^i(x^i \mid s^i) > 0$ and $E_{\bar{Q}^i_{\mu^i}(\cdot|s^i,x^i)}\left[\pi^i(x^i, Y^i)\right] + 4\varepsilon \leq E_{\bar{Q}^i_{\mu^i}(\cdot|s^i,\hat{x}^i)}\left[\pi^i(\hat{x}^i, Y^i)\right]$. By continuity of $\mu^i \mapsto \bar{Q}^i_{\mu^i}$ and the fact that $\lim_{n\to\infty}\mu_n^i = \mu^i$, $\exists n_1$ such that, $\forall n \geq n_1$, $E_{\bar{Q}^i_{\mu_n^i}(\cdot|s^i,x^i)}\left[\pi^i(x^i, Y^i)\right] + 2\varepsilon \leq E_{\bar{Q}^i_{\mu_n^i}(\cdot|s^i,\hat{x}^i)}\left[\pi^i(\hat{x}^i, Y^i)\right]$. It then follows from (10) and $\lim_{n\to\infty} P_{\xi,n}\left(\|\xi_n^i\| \geq \varepsilon\right) = 0$ that $\lim_{n\to\infty}\sigma^i(\mu_n^i, P_{\xi,n})(x^i|s^i) = 0$. But this contradicts $\lim_{n\to\infty}\sigma^i(\mu_n^i, P_{\xi,n})(x^i|s^i) = \sigma^i(x^i \mid s^i) > 0$. □

Proof of Proposition 3. In the next paragraph, we prove the following result: For all σ and $\bar{\theta}^i_\sigma \in \Theta^i(\sigma)$, (a) $Q^i_{\Omega,\bar{\theta}^i_\sigma}(\omega' \mid s^i) = p_{\Omega|S^i}(\omega' \mid s^i)$ for all $s^i \in \mathcal{S}^i, \omega' \in \Omega$ and (b) $Q^i_{\mathbb{X}^{-i},\bar{\theta}^i_\sigma}(x^{-i} \mid \alpha^i) = \sum_{\omega''\in\Omega} p_{\Omega|\mathcal{A}^i}(\omega'' \mid \alpha^i)\prod_{j\neq i}\sigma^j(x^j \mid s^j(\omega''))$ for all $\alpha^i \in \mathcal{A}^i, x^{-i} \in \mathbb{X}^{-i}$. Equivalence between Berk-Nash and ABEE follows immediately from (a), (b), and the fact that expected utility of player i with signal s^i and beliefs $\bar{\theta}_\sigma$

[43] This last claim follows since the weak topology is induced by a family of semi-norms of the form: $\rho(\mu, \mu') = |E_\mu[f] - E_{\mu'}[f]|$ for f continuous and bounded for any μ and μ' in $\Delta(\Theta^i)$.

equals $\sum_{\omega' \in \Omega} Q^i_{\Omega,\bar{\theta}^i_\sigma}(\omega' \mid s^i) \sum_{x^{-i} \in \mathbb{X}^{-i}} Q^i_{\mathbb{X}^{-i},\bar{\theta}^i_\sigma}(x^{-i} \mid \alpha^i(\omega'))\pi^i(\bar{x}^i, x^{-i}, \omega')$.

Proof of (a) and (b): $-K^i(\sigma, \theta^i)$ equals, up to a constant,

$$\sum_{s^i, \tilde{\omega}, \tilde{x}^{-i}} \ln\left(Q^i_{\Omega,\theta^i}(\tilde{\omega} \mid s^i) Q^i_{\mathbb{X}^{-i},\theta^i}(\tilde{x}^{-i} \mid \alpha^i(\tilde{\omega}))\right) \prod_{j \neq i} \sigma^j(\tilde{x}^j \mid s^j(\tilde{\omega})) p_{\Omega|S^i}(\tilde{\omega} \mid s^i) p_{S^i}(s^i)$$
$$= \sum_{s^i, \tilde{\omega}} \ln\left(Q^i_{\Omega,\theta^i}(\tilde{\omega} \mid s^i)\right) p_{\Omega|S^i}(\tilde{\omega} \mid s^i) p_{S^i}(s^i) + \sum_{\tilde{x}^{-i}, \alpha^i \in \mathcal{A}^i} \ln\left(Q^i_{\mathbb{X}^{-i},\theta^i}(\tilde{x}^{-i} \mid \alpha^i)\right) \sum_{\tilde{\omega} \in \alpha^i} \prod_{j \neq i} \sigma^j(\tilde{x}^j \mid s^j(\tilde{\omega})) p_\Omega(\tilde{\omega}).$$

It is straightforward to check that any parameter value that maximizes the above expression satisfies (a) and (b). □

Proof of Lemma 2. The proof uses Claim B, which is stated and proven after this proof. It is sufficient to establish that $\lim_{t\to\infty} \int_{\Theta^i} d^i(\sigma, \theta^i)\mu^i_t(d\theta^i) = 0$ a.s. in $\mathcal{H}$, where $d^i(\sigma, \theta^i) = \inf_{\hat{\theta}^i \in \Theta^i(\sigma)} \|\theta^i - \hat{\theta}^i\|$. Fix $i \in I$ and $h \in \mathbb{H}$. Then, by Bayes' rule,

$$\int_{\Theta^i} d^i(\sigma, \theta^i)\mu^i_t(d\theta^i) = \frac{\int_{\Theta^i} d^i(\sigma, \theta^i) \prod_{\tau=0}^{t-1} \frac{Q^i_{\theta^i}(y^i_\tau | s^i_\tau, x^i_\tau)}{Q^i_{\sigma_\tau}(y^i_\tau | s^i_\tau, x^i_\tau)} \mu^i_0(d\theta^i)}{\int_{\Theta^i} \prod_{\tau=0}^{t-1} \frac{Q^i_{\theta^i}(y^i_\tau | s^i_\tau, x^i_\tau)}{Q^i_{\sigma_\tau}(y^i_\tau | s^i_\tau, x^i_\tau)} \mu^i_0(d\theta^i)} = \frac{\int_{\Theta^i} d^i(\sigma, \theta^i) e^{tK^i_t(h,\theta^i)} \mu^i_0(d\theta^i)}{\int_{\Theta^i} e^{tK^i_t(h,\theta^i)} \mu^i_0(d\theta^i)},$$

where the first equality is well-defined by Assumption 1, full support of μ^i_0, and the fact that $\mathbf{P}^{\mu_0,\phi}(\mathcal{H}) > 0$ implies that all the $Q^i_{\sigma_\tau}$ terms are positive, and where we define $K^i_t(h, \theta^i) = -\frac{1}{t}\sum_{\tau=0}^{t-1} \ln \frac{Q^i_{\sigma_\tau}(y^i_\tau | s^i_\tau, x^i_\tau)}{Q^i_{\theta^i}(y^i_\tau | s^i_\tau, x^i_\tau)}$ for the second equality.[44] For any $\alpha > 0$, define $\Theta^i_\alpha(\sigma) \equiv \{\theta^i \in \Theta^i : d^i(\sigma, \theta^i) < \alpha\}$. Then, for all $\varepsilon > 0$ and $\eta > 0$, $\int_{\Theta^i} d^i(\sigma, \theta^i)\mu^i_t(d\theta^i) \leq \varepsilon + C\frac{A^i_t(h,\sigma,\varepsilon)}{B^i_t(h,\sigma,\eta)}$, where $C \equiv \sup_{\theta^i_1, \theta^i_2 \in \Theta^i} \|\theta^i_1 - \theta^i_2\| < \infty$ (because Θ^i is bounded) and where $A^i_t(h, \sigma, \varepsilon) = \int_{\Theta^i \setminus \Theta^i_\varepsilon(\sigma)} e^{tK^i_t(h,\theta^i)} \mu^i_0(d\theta^i)$ and $B^i_t(h, \sigma, \eta) = \int_{\Theta^i_\eta(\sigma)} e^{tK^i_t(h,\theta^i)} \mu^i_0(d\theta^i)$.The proof concludes by showing that, for all (sufficiently small) $\varepsilon > 0$, $\exists \eta_\varepsilon > 0$ such that $\lim_{t\to\infty} A^i_t(h, \sigma, \varepsilon)/B^i_t(h, \sigma, \eta_\varepsilon) = 0$. This result is achieved in several steps. First, $\forall \varepsilon > 0$, define $K^i_\varepsilon(\sigma) = \inf\{K^i(\sigma, \theta^i) \mid \theta^i \in \Theta^i \setminus \Theta^i_\varepsilon(\sigma)\}$ and $\alpha_\varepsilon = (K^i_\varepsilon(\sigma) - K^i_0(\sigma))/3$, where $K^i_0(\sigma) = \inf_{\theta^i \in \Theta^i} K^i(\sigma, \theta^i)$. By continuity of $K^i(\sigma, \cdot)$, there exist $\bar{\varepsilon}$ and $\bar{\alpha}$ such that, $\forall \varepsilon \leq \bar{\varepsilon}$, $0 < \alpha_\varepsilon \leq \bar{\alpha} < \infty$. Henceforth, let $\varepsilon \leq \bar{\varepsilon}$. It follows that

$$K^i(\sigma, \theta^i) \geq K^i_\varepsilon(\sigma) > K^i_0(\sigma) + 2\alpha_\varepsilon \tag{11}$$

$\forall \theta^i$ such that $d^i(\sigma, \theta^i) \geq \varepsilon$. Also, by continuity of $K^i(\sigma, \cdot)$, $\exists \eta_\varepsilon > 0$ such that, $\forall \theta^i \in$

[44] If, for some θ^i, $Q^i_{\theta^i}(y^i_\tau \mid s^i_\tau, x^i_\tau) = 0$ for some $\tau \in \{0, ..., t-1\}$, then we define $K^i_t(h, \theta^i) = -\infty$ and $\exp\{tK^i_t(h, \theta^i)\} = 0$.

$\Theta^i_{\eta_\varepsilon}$,

$$K^i(\sigma,\theta^i) < K^i_0(\sigma) + \alpha_\varepsilon/2. \tag{12}$$

Second, let $\hat{\Theta}^i = \left\{\theta^i \in \Theta^i : Q^i_{\theta^i}(y^i_\tau \mid s^i_\tau, x^i_\tau) > 0 \text{ for all } \tau\right\}$ and $\hat{\Theta}^i_{\eta_\varepsilon}(\sigma) = \hat{\Theta}^i \cap \Theta^i_{\eta_\varepsilon}(\sigma)$. We now show that $\mu^i_0(\hat{\Theta}^i_{\eta_\varepsilon}(\sigma)) > 0$. By Lemma 1, $\Theta^i(\sigma)$ is nonempty. Pick any $\theta^i \in \Theta^i(\sigma)$. By Assumption 1, $\exists(\theta^i_n)_n$ in Θ^i such that $\lim_{n\to\infty}\theta^i_n = \theta^i$ and $Q^i_{\theta^i_n}(y^i \mid s^i, x^i) > 0$ $\forall y^i \in f^i(\Omega, x^i, \mathbb{X}^{-i})$ and all $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$. In particular, $\exists\theta^i_{\bar{n}}$ such that $d^i(\sigma,\theta^i_{\bar{n}}) < .5\eta_\varepsilon$ and, by continuity of $Q.$, there exists an open set U around $\theta^i_{\bar{n}}$ such that $U \subseteq \hat{\Theta}^i_{\eta_\varepsilon}(\sigma)$. By full support, $\mu^i_0(\hat{\Theta}^i_{\eta_\varepsilon}(\sigma)) > 0$. Next, note that,

$$\begin{aligned}\liminf_{t\to\infty} B^i_t(h,\sigma,\eta_\varepsilon)e^{t(K^i_0(\sigma)+\frac{\alpha_\varepsilon}{2})} &\geq \liminf_{t\to\infty}\int_{\hat{\Theta}^i_{\eta_\varepsilon}(\sigma)} e^{t(K^i_0(\sigma)+\frac{\alpha_\varepsilon}{2}+K^i_t(h,\theta^i))}\mu^i_0(d\theta^i)\\ &\geq \int_{\hat{\Theta}^i_{\eta_\varepsilon}(\sigma)} e^{\lim_{t\to\infty} t(K^i_0(\sigma)+\frac{\alpha_\varepsilon}{2}-K^i(\sigma,\theta^i))}\mu^i_0(d\theta^i) = \infty \end{aligned}\tag{13}$$

a.s. in $\mathcal{H}$, where the first inequality follows because $\hat{\Theta}^i_{\eta_\varepsilon}(\sigma) \subseteq \Theta^i_{\eta_\varepsilon}(\sigma)$ and exp is a positive function, the second inequality follows from Fatou's Lemma and a LLN for non-iid random variables that implies $\lim_{t\to\infty} K^i_t(h,\theta^i) = -K^i(\sigma,\theta^i)$ $\forall\theta^i \in \hat{\Theta}^i$, a.s. in $\mathcal{H}$ (see Claim B(i) below), and the last equality follows from (12) and the fact that $\mu^i_0(\Theta^i_{\eta_\varepsilon}(\sigma)) > 0$.

Next, we consider the term $A^i_t(h,\sigma,\varepsilon)$. Claims B(ii) and B(iii) (see below) imply that $\exists T$ such that, $\forall t \geq T$, $K^i_t(h,\theta^i) < -(K^i_0(\sigma) + (3/2)\alpha_\varepsilon)$ $\forall\theta^i \in \Theta^i\backslash\Theta^i_\varepsilon(\sigma)$, a.s. in $\mathcal{H}$. Thus,

$$\begin{aligned}\lim_{t\to\infty} A^i_t(h,\sigma,\varepsilon)e^{t(K^i_0(\sigma)+\alpha_\varepsilon)} &= \lim_{t\to\infty}\int_{\Theta^i\backslash\Theta^i_\varepsilon(\sigma)} e^{t(K^i_0(\sigma)+\alpha_\varepsilon+K^i_t(h,\theta^i))}\mu^i_0(d\theta^i)\\ &\leq \mu^i_0(\Theta^i\backslash\Theta^i_\varepsilon(\sigma))\lim_{t\to\infty} e^{-t\alpha_\varepsilon/2} = 0.\end{aligned}$$

a.s. in $\mathcal{H}$. The above expression and equation (13) imply that $\lim_{t\to\infty} A^i_t(h,\sigma,\varepsilon)/B^i_t(h,\sigma,\eta_\varepsilon) = 0$ a.s.-$\mathbf{P}^{\mu_0,\phi}$. □

We state and prove Claim B used in the proof above. For any $\xi > 0$, define $\Theta^i_{\sigma,\xi}$ to be the set such that $\theta^i \in \Theta^i_{\sigma,\xi}$ if and only if $Q^i_{\theta^i}(y^i \mid s^i, x^i) \geq \xi$ for all (s^i, x^i, y^i) such that $Q^i_\sigma(y^i \mid s^i, x^i)\sigma^i(x^i \mid s^i)p_{S^i}(s^i) > 0$.

Claim B. For all $i \in I$: **(i)** For all $\theta^i \in \hat{\Theta}^i$, $\lim_{t\to\infty} K^i_t(h,\theta^i) = -K^i(\sigma,\theta^i)$, *a.s. in* $\mathcal{H}$; **(ii)** *There exist* $\xi^* > 0$ *and* T_{ξ^*} *such that,* $\forall t \geq T_{\xi^*}$, $K^i_t(h,\theta^i) < -(K^i_0(\sigma)+(3/2)\alpha_\varepsilon)$ $\forall\theta^i \notin \Theta^i_{\sigma,\xi}$, *a.s. in* $\mathcal{H}$; **(iii)** *For all* $\xi > 0$, $\exists\hat{T}_\xi$ *such that,* $\forall t \geq \hat{T}_\xi$, $K^i_t(h,\theta^i) <$

$-(K_0^i(\sigma) + (3/2)\alpha_\varepsilon)\ \forall \theta^i \in \Theta^i_{\sigma,\xi} \backslash \Theta^i_\varepsilon(\sigma)$, *a.s. in* $\mathcal{H}$.

Proof: Define $freq_t^i(z^i) = \frac{1}{t}\sum_{\tau=0}^{t-1} \mathbf{1}_{z^i}(z_\tau^i)\ \forall z^i \in \mathbb{Z}^i$. K_t^i can be written as $K_t^i(h,\theta^i) = \kappa_{1t}^i(h) + \kappa_{2t}^i(h) + \kappa_{3t}^i(h,\theta^i)$, where $\kappa_{1t}^i(h) = -t^{-1}\sum_{\tau=0}^{t-1}\sum_{z^i\in\mathbb{Z}^i}\left(\mathbf{1}_{z^i}(z_\tau^i) - \bar{P}^i_{\sigma_\tau}(z^i)\right)\ln Q^i_{\sigma_\tau}(z^i)$, $\kappa_{2t}^i(h) = -t^{-1}\sum_{\tau=0}^{t-1}\sum_{z^i\in\mathbb{Z}^i}\bar{P}^i_{\sigma_\tau}(z^i)\ln Q^i_{\sigma_\tau}(z^i)$, and $\kappa_{3t}^i(h,\theta^i) = \sum_{z^i\in\mathbb{Z}^i} freq_t^i(z^i)\ln Q^i_{\theta^i}(z^i)$. The statements made below hold almost surely in $\mathcal{H}$, but we omit this qualification.

First, we show $\lim_{t\to\infty}\kappa_{1t}^i(h) = 0$. Define $l_\tau^i(h,z^i) = \left(\mathbf{1}_{z^i}(z_\tau^i) - \bar{P}^i_{\sigma_\tau}(z^i)\right)\ln Q^i_{\sigma_\tau}(z^i)$ and $L_t^i(h,z^i) = \sum_{\tau=1}^t \tau^{-1} l_\tau^i(h,z^i)\ \forall z^i \in \mathbb{Z}^i$. Fix any $z^i \in \mathbb{Z}^i$. We show that $L_t^i(\cdot,z^i)$ converges a.s. to an integrable, and, therefore, finite function $L_\infty^i(\cdot,z^i)$. To show this, we use martingale convergence results. Let h^t denote the partial history until time t. Since $E_{\mathbf{P}^{\mu_0,\phi}(\cdot|h^t)}\left[l_{t+1}^i(h,z^i)\right] = 0$, then $E_{\mathbf{P}^{\mu_0,\phi}(\cdot|h^t)}\left[L_{t+1}^i(h,z^i)\right] = L_t^i(h,z^i)$ and so $(L_t^i(h,z^i))_t$ is a martingale with respect to $\mathbf{P}^{\mu_0,\phi}$. Next, we show that $\sup_t E_{\mathbf{P}^{\mu_0,\phi}}\left[|L_t^i(h,z^i)|\right] \leq M$ for $M < \infty$. Note that $E_{\mathbf{P}^{\mu_0,\phi}}\left[\left(L_t^i(h,z^i)\right)^2\right] = E_{\mathbf{P}^{\mu_0,\phi}}\left[\sum_{\tau=1}^t \tau^{-2}\left(l_\tau^i(h,z^i)\right)^2 + 2\sum_{\tau'>\tau}\frac{1}{\tau'\tau} l_\tau^i(h,z^i) l_{\tau'}^i(h,z^i)\right]$. Since $(l_t^i)_t$ is a martingale difference sequence, for $\tau' > \tau$, $E_{\mathbf{P}^{\mu_0,\phi}}\left[l_\tau^i(h,z^i) l_{\tau'}^i(h,z^i)\right] = 0$. Therefore, $E_{\mathbf{P}^{\mu_0,\phi}}\left[\left(L_t^i(h,z^i)\right)^2\right] = \sum_{\tau=1}^t \tau^{-2} E_{\mathbf{P}^{\mu_0,\phi}}\left[\left(l_\tau^i(h,z^i)\right)^2\right]$. Note also that $E_{\mathbf{P}^{\mu_0,\phi}(\cdot|h^{\tau-1})}\left[\left(l_\tau^i(h,z^i)\right)^2\right] \leq \left(\ln Q^i_{\sigma_\tau}(z^i)\right)^2 Q^i_{\sigma_\tau}(z^i)$. Therefore, by the law of iterated expectations, $E_{\mathbf{P}^{\mu_0,\phi}}\left[\left(L_t^i(h,z^i)\right)^2\right] \leq \sum_{\tau=1}^t \tau^{-2} E_{\mathbf{P}^{\mu_0,\phi}}\left[\left(\ln Q^i_{\sigma_\tau}(z^i)\right)^2 Q^i_{\sigma_\tau}(z^i)\right]$, which in turn is bounded above by 1 because $(\ln x)^2 x \leq 1\ \forall x \in [0,1]$, where we use the convention that $(\ln 0)^2 0 = 0$. Therefore, $\sup_t E_{\mathbf{P}^{\mu_0,\phi}}\left[|L_t^i|\right] \leq 1$. By Theorem 5.2.8 in Durrett (2010), $L_t^i(h,z^i)$ converges a.s.-$\mathbf{P}^{\mu_0,\phi}$ to a finite $L_\infty^i(h,z^i)$. Thus, by Kronecker's lemma (Pollard (2001), page 105)[45], $\lim_{t\to\infty}\sum_{z^i\in\mathbb{Z}^i}\left\{t^{-1}\sum_{\tau=1}^t\left(1_{z^i}(z_\tau^i) - \bar{P}^i_{\sigma_\tau}(z^i)\right)\ln Q^i_{\sigma_\tau}(z^i)\right\} = 0$. Therefore, $\lim_{t\to\infty}\kappa_{1t}^i(h) = 0$.

Next, consider $\kappa_{2t}^i(h)$. The assumption that $\lim_{t\to\infty}\sigma_t = \sigma$ and continuity of $Q_\sigma^i \ln Q_\sigma^i$ in σ imply that $\lim_{t\to\infty}\kappa_{2t}^i(h) = -\sum_{(s^i,x^i)\in\mathbb{S}^i\times\mathbb{X}^i} E_{Q_\sigma(\cdot|s^i,x^i)}\left[\ln Q_\sigma^i(Y^i \mid s^i,x^i)\right]\sigma^i(x^i \mid s^i)p_{S^i}(s^i)$.

The limits of $\kappa_{1t}^i, \kappa_{2t}^i$ imply that, $\forall\gamma > 0$, $\exists\hat{t}_\gamma$ such that, $\forall t \geq \hat{t}_\gamma$,

$$\left|\kappa_{1t}^i(h) + \kappa_{2t}^i(h) + \sum_{(s^i,x^i)\in\mathbb{S}^i\times\mathbb{X}^i} E_{Q_\sigma(\cdot|s^i,x^i)}\left[\ln Q_\sigma^i(Y^i \mid s^i,x^i)\right]\sigma^i(x^i \mid s^i)p_{S^i}(s^i)\right| \leq \gamma. \quad (14)$$

We now prove (i)-(iii) by characterizing the limit of $\kappa_{3t}^i(h,\theta^i)$.

(i) For all $z^i \in \mathbb{Z}^i$, $\left|freq_t(z^i) - \bar{P}_\sigma^i(z^i)\right| \leq \left|t^{-1}\sum_{\tau=0}^{t-1}\left(\mathbf{1}_{z^i}(z_\tau^i) - \bar{P}^i_{\sigma_\tau}(z^i)\right)\right| + \left|t^{-1}\sum_{\tau=0}^{t-1}\left(\bar{P}^i_{\sigma_\tau}(z^i) - \bar{P}_\sigma^i(z^i)\right)\right|$. The first term in the RHS goes to 0 (the proof is essentially identical to the

[45]This lemma implies that for a sequence $(\ell_t)_t$ if $\sum_\tau \ell_\tau < \infty$, then $\sum_{\tau=1}^t \frac{b_\tau}{b_t}\ell_\tau \to 0$ where $(b_t)_t$ is a nondecreasing positive real valued that diverges to ∞. We can apply the lemma with $\ell_t \equiv t^{-1}l_t$ and $b_t = t$.

proof above that κ_1^i goes to 0). The second term goes to 0 because $\lim_{t\to\infty}\sigma_t = \sigma$ and $\bar{P}_\cdot^i$ is continuous. Thus, $\forall\zeta > 0$, $\exists\hat{t}_\zeta$ such that, $\forall t \geq \hat{t}_\zeta$, $\forall z^i \in \mathbb{Z}^i$,

$$\left|freq_t^i(z^i) - \bar{P}_\sigma^i(z^i)\right| < \zeta \tag{15}$$

Thus, since $\theta^i \in \hat{\Theta}^i$, $\lim_{t\to\infty}\kappa_{3t}^i(h,\theta^i) = \sum_{(s^i,x^i)\in\mathbb{S}^i\times\mathbb{X}^i} E_{Q_\sigma(\cdot|s^i,x^i)}\left[\ln Q_{\theta^i}^i(Y^i \mid s^i,x^i)\right]\sigma^i(x^i \mid s^i)p_{S^i}(s^i)$. This expression and (14) establish part (i).

(ii) For all $\theta^i \notin \Theta_{\sigma,\xi}^i$, let $z_{\theta^i}^i$ be such that $\bar{P}_\sigma^i(z_{\theta^i}^i) > 0$ and $Q_{\theta^i}^i(z_{\theta^i}^i) < \xi$. By (15), $\exists t_{p_L^i/2}$ such $\forall t \geq t_{p_L^i/2}$, $\kappa_{3t}^i(h,\theta^i) \leq freq_t^i(z_{\theta^i}^i)\ln Q_{\theta^i}^i(z_{\theta^i}^i) \leq (p_L^i/2)\ln\xi$ $\forall\theta^i \notin \Theta_{\sigma,\xi}^i$, where $p_L^i = \min_{\mathbb{Z}^i}\{\bar{P}_\sigma^i(z^i) : \bar{P}_\sigma^i(z^i) > 0\}$. This result and (14) imply that, $\forall t \geq t_1 \equiv \max\{t_{p_L^i/2}, \hat{t}_1\}$,

$$\begin{aligned} K_t^i(h,\theta^i) &\leq -\sum_{(s^i,x^i)\in\mathbb{S}^i\times\mathbb{X}^i} E_{Q_\sigma(\cdot|s^i,x^i)}\left[\ln Q_\sigma^i(Y^i \mid s^i,x^i)\right]\sigma^i(x^i \mid s^i)p_{S^i}(s^i) + 1 + \left(p_L^i/2\right)\ln\xi \\ &\leq \#\mathbb{Z}^i + 1 + \left(q_L^i/2\right)\ln\xi, \end{aligned} \tag{16}$$

$\forall\theta^i \notin \Theta_{\sigma,\xi}^i$, where the second line follows from the facts that $\sigma^i(x^i \mid s^i)p_{S^i}(s^i) \leq 1$ and $x\ln(x) \in [-1,0]$ $\forall x \in [0,1]$. In addition, the fact that $K_0^i(\sigma) < \infty$ and $\alpha_\varepsilon \leq \bar{\alpha} < \infty$ $\forall\varepsilon \leq \bar{\varepsilon}$ implies that the RHS of (16) can be made lower than $-(K_0^i(\sigma) + (3/2)\alpha_\varepsilon)$ for some sufficiently small ξ^*.

(iii) For any $\xi > 0$, let $\zeta_\xi = -\alpha_\varepsilon/(\#\mathbb{Z}^i 4\ln\xi) > 0$. By (15), $\exists\hat{t}_{\zeta_\xi}$ such that, $\forall t \geq \hat{t}_{\zeta_\xi}$,

$$\begin{aligned} \kappa_{3t}^i(h,\theta^i) &\leq \sum_{\{z^i:\bar{P}_\sigma^i(z^i)>0\}} freq_t^i(z^i)\ln Q_{\theta^i}^i(z^i) \leq \sum_{\{z^i:\bar{P}_\sigma^i(z^i)>0\}} \left(\bar{P}_\sigma^i(z^i) - \zeta_\xi\right)\ln Q_{\theta^i}^i(z^i) \\ &\leq \sum_{(s^i,x^i)\in\mathbb{S}^i\times\mathbb{X}^i} E_{Q_\sigma(\cdot|s^i,x^i)}\left[\ln Q_{\theta^i}^i(Y^i \mid s^i,x^i)\right]\sigma^i(x^i \mid s^i)p_{S^i}(s^i) - \#\mathbb{Z}^i\zeta_\xi\ln\xi, \end{aligned}$$

$\forall\theta^i \in \Theta_{\sigma,\xi}^i$ (since $Q_{\theta^i}^i(z^i) \geq \xi$ $\forall z^i$ such that $\bar{P}_\sigma^i(z^i) > 0$). The above expression, the fact that $\alpha_\varepsilon/4 = -\#\mathbb{Z}^i\zeta_\xi\ln\xi$, and (14) imply that, $\forall t \geq \hat{T}_\xi \equiv \max\{\hat{t}_{\zeta_\xi}, \hat{t}_{\alpha_\varepsilon/4}\}$, $K_t^i(h,\theta^i) < -K^i(\sigma,\theta^i) + \alpha_\varepsilon/2$ $\forall\theta^i \in \Theta_{\sigma,\xi}^i$. This result and (11) imply the desired result. □

Online Appendix

A Example: Trading with adverse selection

In this section, we provide the formal details for the trading environment in Example 2.5. Let $p \in \Delta(\mathbb{A} \times \mathbb{V})$ be the true distribution; we use subscripts, such as p_A and $p_{V|A}$, to denote the corresponding marginal and conditional distributions. Let $\mathbb{Y} = \mathbb{A} \times \mathbb{V} \cup \{\square\}$ denote the space of observable consequences, where $\square$ will be a convenient way to represent the fact that there is no trade. We denote the random variable taking values in $\mathbb{V} \cup \{\square\}$ by $\hat{V}$. Notice that the state space in this example is $\Omega = \mathbb{A} \times \mathbb{V}$.

Partial feedback is represented by the function $f^P : \mathbb{X} \times \mathbb{A} \times \mathbb{V} \to \mathbb{Y}$ such that $f^P(x, a, v) = (a, v)$ if $a \leq x$ and $f^P(x, a, v) = (a, \square)$ if $a > x$. Full feedback is represented by $f^F(x, a, v) = (a, v)$. In all cases, payoffs are given by $\pi : \mathbb{X} \times \mathbb{Y} \to \mathbb{R}$, where $\pi(x, a, v) = v - x$ if $a \leq x$ and 0 otherwise. The objective distribution for the case of partial feedback, Q^P, is, $\forall x \in \mathbb{X}$, $\forall (a, v) \in \mathbb{A} \times \mathbb{V}$, $Q^P(a, v \mid x) = p(a, v) 1_{\{x \geq a\}}(x)$, and, $\forall x \in \mathbb{X}$, $\forall a \in \mathbb{A}$, $Q^P(a, \square \mid x) = p_A(a) 1_{\{x < a\}}(x)$. The objective distribution for the case of full feedback, Q^F, is, $\forall x \in \mathbb{X}$, $\forall (a, v) \in \mathbb{A} \times \mathbb{V}$, $Q^F(a, v \mid x) = p(a, v)$, and, $\forall x \in \mathbb{X}$, $\forall a \in \mathbb{A}$, $Q^F(a, \square \mid x) = 0$.

The buyer knows the environment except for the distribution $p \in \Delta(\mathbb{A} \times \mathbb{V})$. Then, for any distribution in the subjective model, Q_θ, the perceived expected profit from choosing $x \in \mathbb{X}$ is

$$E_{Q_\theta(\cdot|x)}[\pi(x, A, \hat{V})] = \sum_{(a,v) \in \mathbb{A} \times \mathbb{V}} 1_{\{x \geq a\}}(x) \, (v - x) \, Q_\theta(a, v \mid x). \tag{17}$$

The buyer has either one of two misspecifications over p captured by the parameter sets $\Theta_I = \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$ (i.e., independent beliefs) or $\Theta_A = \times_j \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$ (i.e., analogy-based beliefs) defined in the main text. Thus, combining feedback and parameter sets, we have four cases to consider, and, for each case, we write down the subjective model and wKLD function.

Cursed equilibrium. Feedback is f^F and the parameter set is Θ_I. The subjective model is, $\forall x \in \mathbb{X}$, $\forall (a, v) \in \mathbb{A} \times \mathbb{V}$, $Q^C_\theta(a, v \mid x) = \theta_A(a) \theta_V(v)$, and, $\forall x \in \mathbb{X}$, $\forall a \in \mathbb{A}$, $Q^C_\theta(a, \square \mid x) = 0$, where $\theta = (\theta_A, \theta_V) \in \Theta_I$.[46] This is an analogy-based game. From

[46] In fact, the symbol $\square$ is not necessary for this example, but we keep it so that all feedback functions are defined over the same space of consequences.

(17), the perceived expected profit from $x \in \mathbb{X}$ is

$$Pr_{\theta_A}(A \leq x)\left(E_{\theta_V}[V] - x\right), \tag{18}$$

where Pr_{θ_A} denotes probability with respect to θ_A and E_{θ_V} denotes expectation with respect to θ_V. Also, for all (pure) strategies $x \in \mathbb{X}$, the wKLD function is[47]

$$K^C(x,\theta) = E_{Q^F(\cdot|x)}\big[\ln \frac{Q^F(A,\hat{V} \mid x)}{Q^C_\theta(A,\hat{V} \mid x)}\big] = \sum_{(a,v)\in\mathbb{A}\times\mathbb{V}} p(a,v) \ln \frac{p(a,v)}{\theta_A(a)\theta_V(v)}.$$

For each $x \in \mathbb{X}$, $\theta(x) = (\theta_A(x), \theta_V(x)) \in \Theta_I = \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$, where $\theta_A(x) = p_A$ and $\theta_V(x) = p_V$ is the unique parameter value that minimizes $K^C(x,\cdot)$. Together with (18), we obtain equation Π^{CE} in the main text.

Behavioral equilibrium (naive version). Feedback is f^P and the parameter set is Θ_I. The subjective model is, $\forall x \in \mathbb{X}$, $\forall (a,v) \in \mathbb{A} \times \mathbb{V}$, $Q^{BE}_\theta(a,v \mid x) = \theta_A(a)\theta_V(v)1_{\{x \geq a\}}(x)$, and, $\forall x \in \mathbb{X}$, $\forall a \in \mathbb{A}$, $Q^{BE}_\theta(a,\square \mid x) = \theta_A(a)1_{\{x<a\}}(x)$, where $\theta = (\theta_A, \theta_V) \in \Theta_I$. From (17), perceived expected profit from $x \in \mathbb{X}$ is as in equation (18). Also, for all (pure) strategies $x \in \mathbb{X}$, the wKLD function is

$$\begin{aligned} K^{BE}(x,\theta) &= E_{Q^P(\cdot|x)}\big[\ln \frac{Q^P(A,\hat{V} \mid x)}{Q^{BE}_\theta(A,\hat{V} \mid x)}\big] \\ &= \sum_{\{a\in\mathbb{A}:a>x\}} p_A(a) \ln \frac{p_A(a)}{\theta_A(a)} + \sum_{\{(a,v)\in\mathbb{A}\times\mathbb{V}:a\leq x\}} p(a,v) \ln \frac{p(a,v)}{\theta_A(a)\theta_V(v)}. \end{aligned}$$

For each $x \in \mathbb{X}$, $\theta(x) = (\theta_A(x), \theta_V(x)) \in \Theta_I = \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$, where $\theta_A(x) = p_A$ and $\theta_V(x)(v) = p_{V|A}(v \mid A \leq x)\ \forall v \in \mathbb{V}$ is the unique parameter value that minimizes $K^{BE}(x,\cdot)$. Together with (18), we obtain equation Π^{BE} in the main text.

Analogy-based expectations equilibrium. Feedback is f^F and the parameter set is Θ_A. The subjective model is, $\forall x \in \mathbb{X}$, $\forall (a,v) \in \mathbb{A} \times \mathbb{V}_j$, all $j = 1, ..., k$, $Q^{ABEE}_\theta(a,v \mid x) = \theta_j(a)\theta_V(v)$, and, $\forall x \in \mathbb{X}$, $\forall a \in \mathbb{A}$, $Q^{ABEE}_\theta(a,\square \mid x) = 0$, where $\theta = (\theta_1, ..., \theta_k, \theta_V) \in \Theta_A$. This is an analogy-based game. From (17), perceived expected profit from $x \in \mathbb{X}$ is

$$\sum_{j=1}^{k} Pr_{\theta_V}(V \in \mathbb{V}_j)\left\{Pr_{\theta_j}(A \leq x \mid V \in \mathbb{V}_j)\left(E_{\theta_V}[V \mid V \in \mathbb{V}_j] - x\right)\right\}. \tag{19}$$

[47] In all cases, the extension to mixed strategies is straightforward.

Also, for all (pure) strategies $x \in \mathbb{X}$, the wKLD function is

$$K^{ABEE}(x,\theta) = E_{Q^F(\cdot|x)}\big[\ln \frac{Q^F(A,\hat{V} \mid x)}{Q_\theta^{ABEE}(A,\hat{V} \mid x)}\big] = \sum_{j=1}^{k} \sum_{(a,v)\in \mathbb{A}\times\mathbb{V}_j} p(a,v) \ln \frac{p(a,v)}{\theta_j(a)\theta_V(v)}.$$

For each $x \in \mathbb{X}$, $\theta(x) = (\theta_1(x), ..., \theta_k(x), \theta_V(x)) \in \Theta_A = \times_j \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$, where $\theta_j(x)(a) = p_{A|V_j}(a \mid V \in \mathbb{V}_j)\ \forall a \in \mathbb{A}$ and $\theta_V(x) = p_V$ is the unique parameter value that minimizes $K^{ABEE}(x,\cdot)$. Together with (19), we obtain equation Π^{ABEE} in the main text.

Behavioral equilibrium (naive version) with analogy classes. It is natural to also consider a case, unexplored in the literature, where feedback f^P is partial and the subjective model is parameterized by Θ_A. Suppose that the buyer's behavior has stabilized to some price x^*. Due to the possible correlation across analogy classes, the buyer might now believe that deviating to a different price $x \neq x^*$ affects her valuation. In particular, the buyer might have multiple beliefs at x^*. To obtain a natural equilibrium refinement, we assume that the buyer also observes the analogy class that contains her realized valuation, whether she trades or not, and that $\Pr(V \in \mathbb{V}_j, A \leq x) > 0$ for all $j = 1, ..., k$ and $x \in \mathbb{X}$.[48] We denote this new feedback assumption by a function $f^{P^*} : \mathbb{X} \times \mathbb{A} \times \mathbb{V} \to \mathbb{Y}^*$ where $\mathbb{Y}^* = \mathbb{A} \times \mathbb{V} \cup \{1, ..., k\}$ and $f^{P^*}(x,a,v) = (a,v)$ if $a \leq x$ and $f^{P^*}(x,a,v) = (a,j)$ if $a > x$ and $v \in \mathbb{V}_j$. The objective distribution given this feedback function is, $\forall x \in \mathbb{X}$, $\forall (a,v) \in \mathbb{A} \times \mathbb{V}$, $Q^{P^*}(a,v \mid x) = p(a,v)1_{\{x \geq a\}}(x)$, and, $\forall x \in \mathbb{X}$, $\forall a \in \mathbb{A}$ and all $j = 1, ..., k$, $Q^{P^*}(a,j \mid x) = p_{A|V_j}(a \mid V \in \mathbb{V}_j)p_V(\mathbb{V}_j)1_{\{x<a\}}(x)$. The subjective model is, $\forall x \in \mathbb{X}$, $\forall (a,v) \in \mathbb{A} \times \mathbb{V}_j$ and all $j = 1, ..., k$, $Q_\theta^{BEA}(a,v \mid x) = \theta_j(a)\theta_V(v)1_{\{x \geq a\}}(x)$, and, $\forall x \in \mathbb{X}$, $\forall (a,v) \in \mathbb{A} \times \mathbb{V}_j$ and all $j = 1, ..., k$, $Q_\theta^{BEA}(a,j \mid x) = \theta_j(a)\left(\sum_{v\in\mathbb{V}_j} \theta_V(v)\right) 1_{\{x<a\}}(x)$, where $\theta = (\theta_1, \theta_2, ..., \theta_k, \theta_V) \in \Theta_A$. In particular, from (17), perceived expected profit from $x \in \mathbb{X}$ is as in equation (19).

[48] Alternatively, and more naturally, we could require the equilibrium to be the limit of a sequence of mixed strategy equilibria with the property that all prices are chosen with positive probability.

Also, for all (pure) strategies $x \in \mathbb{X}$, the wKLD function is

$$K^{BEA}(x,\theta) = E_{Q^{P^*}(\cdot|x)}\big[\ln \frac{Q^{P^*}(A,\hat{V} \mid x)}{Q_\theta^{BEA}(A,\hat{V} \mid x)}\big] = \sum_{j=1}^{k} \sum_{\{(a,v)\in \mathbb{A}\times\mathbb{V}_j : a\leq x\}} p(a,v) \ln \frac{p(a,v)}{\theta_j(a)\theta_V(v)}$$
$$+ \sum_{\{(a,j)\in\mathbb{A}\times\{1,\ldots,k\}: a>x\}} p_{A|V_j}(a \mid V\in\mathbb{V}_j) p_V(\mathbb{V}_j) \ln \frac{p_{A|V_j}(a \mid V\in\mathbb{V}_j) p_V(\mathbb{V}_j)}{\theta_j(a)\sum_{v\in\mathbb{V}_j}\theta_V(v)}.$$

For each $x \in \mathbb{X}$, $\theta(x) = (\theta_1(x),...,\theta_k(x),\theta_V(x)) \in \Theta_A = \times_j \Delta(\mathbb{A}) \times \Delta(\mathbb{V})$, where $\theta_j(x)(a) = p_{A|V_j}(a \mid V \in \mathbb{V}_j)$ $\forall a \in \mathbb{A}$ and $\theta_V(x)(v) = p_{V|A}(v \mid V \in \mathbb{V}_j,\, A \leq x) p_V(\mathbb{V}_j)$ $\forall v \in \mathbb{V}_j$, all $j = 1,...,k$ is the unique parameter value that minimizes $K^{BEA}(x,\cdot)$. Together with (19), we obtain $\Pi^{BEA}(x,x^*) = \sum_{i=j}^{k} \Pr(V \in \mathbb{V}_j) \Pr(A \leq x \mid V \in \mathbb{V}_j)\big(E\big[V \mid V \in \mathbb{V}_j, A \leq x^*\big] - x\big)$.

B Proof of converse result: Theorem 3

Let $(\bar{\mu}^i)_{i\in I}$ be a belief profile that supports σ as an equilibrium. Consider the following policy profile $\phi = (\phi_t^i)_{i,t}$: For all $i \in I$ and all t,

$$(\mu^i, s^i, \xi^i) \mapsto \phi_t^i(\mu^i, s^i, \xi^i) \equiv \begin{cases} \varphi^i(\bar{\mu}^i, s^i, \xi^i) & \text{if } \max_{i\in I} ||\bar{Q}^i_{\mu^i} - \bar{Q}^i_{\bar{\mu}^i}|| \leq \frac{1}{2C}\varepsilon_t \\ \varphi^i(\mu^i, s^i, \xi^i) & \text{otherwise,} \end{cases}$$

where φ^i is an arbitrary selection from Ψ^i, $C \equiv \max_I \left\{\#\mathbb{Y}^i \times \sup_{\mathbb{X}^i\times\mathbb{Y}^i} |\pi^i(x^i,y^i)|\right\} < \infty$, and the sequence $(\varepsilon_t)_t$ will be defined below. For all $i \in I$, fix any prior μ_0^i with full support on Θ^i such that $\mu_0^i(\cdot|\Theta^i(\sigma)) = \bar{\mu}^i$ (where for any $A \subset \Theta$ Borel, $\mu(\cdot|A)$ is the conditional probability given A).

We now show that if $\varepsilon_t \geq 0$ $\forall t$ and $\lim_{t\to\infty} \varepsilon_t = 0$, then ϕ is asymptotically optimal. Throughout this argument, we fix an arbitrary $i \in I$. Abusing notation, let $U^i(\mu^i, s^i, \xi^i, x^i) = E_{\bar{Q}_{\mu^i}(\cdot|s^i,x^i)}[\pi^i(x^i, Y^i)] + \xi^i(x^i)$. It suffices to show that

$$U^i(\mu^i, s^i, \xi^i, \phi_t^i(\mu^i, s^i, \xi^i)) \geq U^i(\mu^i, s^i, \xi^i, x^i) - \varepsilon_t \tag{20}$$

for all (i,t), all (μ^i, s^i, ξ^i), and all x^i. By construction of ϕ, equation (20) is satisfied if $\max_{i\in I} ||\bar{Q}^i_{\mu^i} - \bar{Q}^i_{\bar{\mu}^i}|| > \frac{1}{2C}\varepsilon_t$. If, instead, $\max_{i\in I} ||\bar{Q}^i_{\mu^i} - \bar{Q}^i_{\bar{\mu}^i}|| \leq \frac{1}{2C}\varepsilon_t$, then

$$U^i(\bar{\mu}^i, s^i, \xi^i, \phi_t^i(\mu^i, s^i, \xi^i)) = U^i(\bar{\mu}^i, s^i, \xi^i, \varphi^i(\bar{\mu}^i, s^i, \xi^i)) \geq U^i(\bar{\mu}^i, s^i, \xi^i, x^i), \tag{21}$$

$\forall x^i \in \mathbb{X}^i$. Moreover, $\forall x^i$,

$$\begin{aligned}\left|U^i(\bar{\mu}^i, s^i, \xi^i, x^i) - U^i(\mu^i, s^i, \xi^i, x^i)\right| &= \Big|\sum_{y^i \in \mathbb{Y}^i} \pi(x^i, y^i)\big(\bar{Q}^i_{\bar{\mu}^i}(y^i \mid s^i, x^i) - \bar{Q}^i_{\mu^i}(y^i \mid s^i, x^i)\big)\Big| \\ &\leq \sup_{\mathbb{X}^i \times \mathbb{Y}^i} |\pi^i(x^i, y^i)| \sum_{y^i \in \mathbb{Y}^i} \left|\big(\bar{Q}^i_{\bar{\mu}^i}(y^i \mid s^i, x^i) - \bar{Q}^i_{\mu^i}(y^i \mid s^i, x^i)\big)\right| \\ &\leq \sup_{\mathbb{X}^i \times \mathbb{Y}^i} |\pi^i(x^i, y^i)| \times \#\mathbb{Y}^i \times \max_{y^i, x^i, s^i} \left|\bar{Q}^i_{\bar{\mu}^i}(y^i \mid s^i, x^i) - \bar{Q}^i_{\mu^i}(y^i \mid s^i, x^i)\right|\end{aligned}$$

so by our choice of C, $|U^i(\bar{\mu}^i, s^i, \xi^i, x^i) - U^i(\mu^i, s^i, \xi^i, x^i)| \leq 0.5\varepsilon_t \ \forall x^i$. Therefore, equation (21) implies equation (20); thus ϕ is asymptotically optimal if $\varepsilon_t \geq 0 \ \forall t$ and $\lim_{t\to\infty} \varepsilon_t = 0$.

We now construct a sequence $(\varepsilon_t)_t$ such that $\varepsilon_t \geq 0 \ \forall t$ and $\lim_{t\to\infty} \varepsilon_t = 0$. Let $\bar{\phi}^i = (\bar{\phi}^i_t)_t$ be such that $\bar{\phi}^i_t(\mu^i, \cdot, \cdot) = \varphi^i(\bar{\mu}^i, \cdot, \cdot) \ \forall \mu^i$; i.e., $\bar{\phi}^i$ is a stationary policy that maximizes utility under the assumption that the belief is always $\bar{\mu}^i$. Let $\zeta^i(\mu^i) \equiv 2C||\bar{Q}^i_{\mu^i} - \bar{Q}^i_{\bar{\mu}^i}||$ and suppose (the proof is at the end) that

$$\boldsymbol{P}^{\mu_0, \bar{\phi}}(\lim_{t\to\infty} \max_{i\in I} |\zeta^i(\mu^i_t(h))| = 0) = 1 \tag{22}$$

(recall that $\boldsymbol{P}^{\mu_0,\bar{\phi}}$ is the probability measure over $\mathbb{H}$ induced by the policy profile $\bar{\phi}$; by definition of $\bar{\phi}$, $\boldsymbol{P}^{\mu_0,\bar{\phi}}$ does not depend on μ_0). Then by the 2nd Borel-Cantelli lemma (Billingsley (1995), pages 59-60), for any $\gamma > 0$, $\sum_t \boldsymbol{P}^{\mu_0,\bar{\phi}}(\max_{i\in I} |\zeta^i(\mu^i_t(h))| \geq \gamma) < \infty$. Hence, for any $a > 0$, there exists a sequence $(\tau(j))_j$ such that

$$\sum_{t \geq \tau(j)} \boldsymbol{P}^{\mu_0,\bar{\phi}}\big(\max_{i\in I} |\zeta^i(\mu^i_t(h))| \geq 1/j\big) < \frac{3}{a} 4^{-j} \tag{23}$$

and $\lim_{j\to\infty} \tau(j) = \infty$. For all $t \leq \tau(1)$, we set $\varepsilon_t = 3C$, and, for any $t > \tau(1)$, we set $\varepsilon_t \equiv 1/N(t)$, where $N(t) \equiv \sum_{j=1}^{\infty} 1\{\tau(j) \leq t\}$. Observe that, since $\lim_{j\to\infty} \tau(j) = \infty$, $N(t) \to \infty$ as $t \to \infty$ and thus $\varepsilon_t \to 0$.

Next, we show that $\mathbf{P}^{\mu_0,\phi}(\lim_{t\to\infty} \|\sigma_t(h^\infty) - \sigma\| = 0) = 1$, where $(\sigma_t)_t$ is the sequence of intended strategies given ϕ, i.e., $\sigma^i_t(h)(x^i \mid s^i) = P_\xi(\xi^i : \phi^i_t(\mu^i_t(h), s^i, \xi^i) = x^i)$. Observe that, by definition, $\sigma^i(x^i \mid s^i) = P_\xi\big(\xi^i : x^i \in \arg\max_{\hat{x}^i \in \mathbb{X}^i} E_{\bar{Q}_{\bar{\mu}^i}(\cdot|s^i,\hat{x}^i)}[\pi^i(\hat{x}^i, Y^i)] + \xi^i(\hat{x}^i)\big)$. Since $\varphi^i \in \Psi^i$, it follows that we can write $\sigma^i(x^i \mid s^i) = P_\xi(\xi^i : \varphi^i(\bar{\mu}^i, s^i, \xi^i) = x^i)$. Let $H \equiv \{h\colon \ \|\sigma_t(h) - \sigma\| = 0, \text{ for all } t\}$. It is sufficient to show that $\mathbf{P}^{\mu_0,\phi}(H) = 1$.

To show this, observe that

$$\begin{aligned}
\mathbf{P}^{\mu_0,\phi}(H) \geq & \mathbf{P}^{\mu_0,\phi}\left(\cap_t\{\max_i \zeta^i(\mu_t) \leq \varepsilon_t\}\right) \\
= & \prod_{t=\tau(1)+1}^{\infty} \mathbf{P}^{\mu_0,\phi}\left(\max_i \zeta^i(\mu_t) \leq \varepsilon_t \mid \cap_{l<t}\{\max_i \zeta^i(\mu_l) \leq \varepsilon_l\}\right) \\
= & \prod_{t=\tau(1)+1}^{\infty} \mathbf{P}^{\mu_0,\bar{\phi}}\left(\max_i \zeta^i(\mu_t) \leq \varepsilon_t \mid \cap_{l<t}\{\max_i \zeta^i(\mu_l) \leq \varepsilon_l\}\right) \\
= & \mathbf{P}^{\mu_0,\bar{\phi}}\left(\cap_{t>\tau(1)}\{\max_i \zeta^i(\mu_t) \leq \varepsilon_t\}\right),
\end{aligned}$$

where the second line omits the term $\mathbf{P}^{\mu_0,\phi}(\max_i \zeta^i(\mu_t) < \varepsilon_t \text{ for all } t \leq \tau(1))$ because it is equal to 1 (since $\varepsilon_t \geq 3C \ \forall t \leq \tau(1)$); the third line follows from the fact that $\phi^i_{t-1} = \bar{\phi}^i_{t-1}$ if $\zeta^i(\mu_{t-1}) \leq \varepsilon_{t-1}$, so the probability measure is equivalently given by $\boldsymbol{P}^{\mu_0,\bar{\phi}}$; and where the last line also uses the fact that $\boldsymbol{P}^{\mu_0,\bar{\phi}}(\max_i \zeta^i(\mu_t) < \varepsilon_t \text{ for all } t \leq \tau(1)) = 1$. In addition, $\forall a > 0$,

$$\begin{aligned}
\boldsymbol{P}^{\mu_0,\bar{\phi}}\left(\cap_{t>\tau(1)}\{\max_i \zeta^i(\mu_t) \leq \varepsilon_t\}\right) = & \boldsymbol{P}^{\mu_0,\bar{\phi}}\left(\cap_{n\in\{1,2,\ldots\}} \cap_{\{t>\tau(1):N(t)=n\}} \{\max_i \zeta^i(\mu_t) \leq n^{-1}\}\right) \\
\geq & 1 - \sum_{n=1}^{\infty} \sum_{\{t:N(t)=n\}} \boldsymbol{P}^{\mu_0,\bar{\phi}}\left(\max_i \zeta^i(\mu_t) \geq n^{-1}\right) \\
\geq & 1 - \sum_{n=1}^{\infty} \frac{3}{a} 4^{-n} = 1 - \frac{1}{a},
\end{aligned}$$

where the last line follows from (23). Thus, we have shown that $\mathbf{P}^{\mu_0,\phi}(H) \geq 1 - 1/a$ $\forall a > 0$; hence, $\mathbf{P}^{\mu_0,\phi}(H) = 1$.

We conclude the proof by showing that equation (22) indeed holds. Observe that σ is trivially stable under $\bar{\phi}$. By Lemma 2, $\forall i \in I$ and all open sets $U^i \supseteq \Theta^i(\sigma)$,

$$\lim_{t\to\infty} \mu^i_t\left(U^i\right) = 1 \tag{24}$$

$a.s. - \boldsymbol{P}^{\mu_0,\bar{\phi}}$ (over $\mathbb{H}$). Let $\mathcal{H}$ denote the set of histories such that $x^i_t(h) = x^i$ and $s^i_t(h) = s^i$ implies that $\sigma^i(x^i \mid s^i) > 0$. By definition of $\bar{\phi}$, $\boldsymbol{P}^{\mu_0,\bar{\phi}}(\mathcal{H}) = 1$. Thus, it suffices to show that $\lim_{t\to\infty} \max_{i\in I} |\zeta^i(\mu^i_t(h))| = 0$ a.s.-$\boldsymbol{P}^{\mu_0,\bar{\phi}}$ over $\mathcal{H}$. To do this, take

any $A \subseteq \Theta$ that is closed. By equation (24), $\forall i \in I$, and almost all $h \in \mathcal{H}$,

$$\limsup_{t\to\infty} \int 1_A(\theta)\mu^i_{t+1}(d\theta) = \limsup_{t\to\infty} \int 1_{A\cap\Theta^i(\sigma)}(\theta)\mu^i_{t+1}(d\theta).$$

Moreover,

$$\begin{aligned}\int 1_{A\cap\Theta^i(\sigma)}(\theta)\mu^i_{t+1}(d\theta) \leq \int 1_{A\cap\Theta^i(\sigma)}(\theta) \left\{ \frac{\prod_{\tau=1}^t Q^i_\theta(y^i_\tau \mid s^i_\tau, x^i_\tau)\mu^i_0(d\theta)}{\int_{\Theta^i(\sigma)} \prod_{\tau=1}^t Q^i_\theta(y^i_\tau \mid s^i_\tau, x^i_\tau)\mu^i_0(d\theta)} \right\} \\ = \mu^i_0(A \mid \Theta^i(\sigma)) = \bar{\mu}^i(A),\end{aligned}$$

where the first inequality follows from the fact that $\Theta^i(\sigma) \subseteq \Theta^i$; the first equality follows from the fact that, since $h \in \mathcal{H}$, the fact that the game is weakly identified given σ implies that $\prod_{\tau=1}^t Q^i_\theta(y^i_\tau \mid s^i_\tau, x^i_\tau)$ is constant with respect to θ $\forall\theta \in \Theta^i(\sigma)$, and the last equality follows from our choice of μ^i_0. Therefore, we established that a.s.-$\boldsymbol{P}^{\mu_0,\bar{\phi}}$ over $\mathcal{H}$, $\limsup_{t\to\infty} \mu^i_{t+1}(h)(A) \leq \bar{\mu}^i(A)$ for A closed. By the portmanteau lemma, this implies that, a.s. -$\boldsymbol{P}^{\mu_0,\bar{\phi}}$ over $\mathcal{H}$, $\lim_{t\to\infty} \int_\Theta f(\theta)\mu^i_{t+1}(h)(d\theta) = \int_\Theta f(\theta)\bar{\mu}^i(d\theta)$ for any f real-valued, bounded and continuous. Since, by assumption, $\theta \mapsto Q^i_\theta(y^i \mid s^i, x^i)$ is bounded and continuous, the previous result applies to $Q^i_\theta(y^i \mid s^i, x^i)$, and since y, s, x take a finite number of values, this result implies that $\lim_{t\to\infty} ||\bar{Q}^i_{\mu^i_t(h)} - \bar{Q}^i_{\bar{\mu}^i}|| = 0$ $\forall i \in I$ a.s. -$\boldsymbol{P}^{\mu_0,\bar{\phi}}$ over $\mathcal{H}$. □

C Non-myopic players

In the main text, we proved the results for the case where players are myopic. Here, we assume that players maximize discounted expected payoffs, where $\delta^i \in [0, 1)$ is the discount factor of player i. In particular, players can be forward looking and decide to experiment. Players believe, however, that they face a stationary environment and, therefore, have no incentives to influence the future behavior of other players. We assume for simplicity that players know the distribution of their own payoff perturbations.

Because players believe that they face a stationary environment, they solve a (subjective) dynamic optimization problem that can be cast recursively as follows. By the Principle of Optimality, $V^i(\mu^i, s^i)$ denotes the maximum expected discounted payoffs (i.e., the value function) of player i who starts a period by observing signal s^i and by

holding belief μ^i if and only if

$$V^i(\mu^i, s^i) = \int_{\Xi^i} \left\{ \max_{x^i \in \mathbb{X}^i} E_{\bar{Q}_{\mu^i}(\cdot|s^i,x^i)} \left[\pi^i(x^i, Y^i) + \xi^i(x^i) + \delta E_{p_{Si}} \left[V^i(\hat{\mu}^i, S^i) \right] \right] \right\} P_\xi(d\xi^i), \tag{25}$$

where $\hat{\mu}^i = B^i(\mu^i, s^i, x^i, Y^i)$ is the updated belief. For all (μ^i, s^i, ξ^i), let

$$\Phi^i(\mu^i, s^i, \xi^i) = \arg\max_{x^i \in \mathbb{X}^i} E_{\bar{Q}_{\mu^i}(\cdot|s^i,x^i)} \left[\pi^i(x^i, Y^i) + \xi^i(x^i) + \delta E_{p_{Si}} \left[V^i(\hat{\mu}^i, S^i) \right] \right].$$

The proof of the next lemma relies on standard arguments and is, therefore, omitted.[49]

Lemma 3. *There exists a unique solution V^i to the Bellman equation (25); this solution is bounded in $\Delta(\Theta^i) \times \mathbb{S}^i$ and continuous as a function of μ^i. Moreover, Φ^i is single-valued and continuous with respect to μ^i, a.s.- P_ξ.*

Because players believe they face a stationary environment with i.i.d. perturbations, it is without loss of generality to restrict behavior to depend on the state of the recursive problem. Optimality of a policy is defined as usual (with the requirement that $\phi_t^i \in \Phi^i \ \forall t$).

Lemma 2 implies that the *support* of posteriors converges, but posteriors need not converge. We can always find, however, a subsequence of posteriors that converges. By continuity of dynamic behavior in beliefs, the stable strategy profile is dynamically optimal (in the sense of solving the dynamic optimization problem) given this convergent posterior. For weakly identified games, the convergent posterior is a fixed point of the Bayesian operator. Thus, the players' limiting strategies will provide no new information. Since the value of experimentation is nonnegative, it follows that the stable strategy profile must also be myopically optimal (in the sense of solving the optimization problem that ignores the future), which is the definition of optimality used in the definition of Berk-Nash equilibrium. Thus, we obtain the following characterization of the set of stable strategy profiles when players follow optimal policies.

Theorem 4. *Suppose that a strategy profile σ is stable under an optimal policy profile for a perturbed and weakly identified game. Then σ is a Berk-Nash equilibrium of the game.*

[49] Doraszelski and Escobar (2010) study a similarly perturbed version of the Bellman equation.

Proof. The first part of the proof is identical to the proof of Theorem 2. Here, we prove that, given that $\lim_{j\to\infty}\sigma_{t(j)}=\sigma$ and $\lim_{j\to\infty}\mu^i_{t(j)}=\mu^i_\infty\in\Delta(\Theta^i(\sigma))\ \forall i$, then, $\forall i$, σ^i is optimal for the perturbed game given $\mu^i_\infty\in\Delta(\Theta^i)$, *i.e.*, $\forall(s^i,x^i)$,

$$\sigma^i(x^i\mid s^i)=P_\xi\left(\xi^i:\psi^i(\mu^i_\infty,s^i,\xi^i)=\{x^i\}\right), \tag{26}$$

where $\psi^i(\mu^i_\infty,s^i,\xi^i)\equiv\arg\max_{x^i\in\mathbb{X}^i}E_{\bar{Q}^i_{\mu^i_\infty}(\cdot|s^i,x^i)}\left[\pi^i(x^i,Y^i)\right]+\xi^i(x^i)$.

To establish (26), fix $i\in I$ and $s^i\in\mathbb{S}^i$. Then

$$\begin{aligned}\lim_{j\to\infty}\sigma^i_{t(j)}(h)(x^i|s^i)&=\lim_{j\to\infty}P_\xi\left(\xi^i:\phi^i_{t(j)}(\mu^i_{t(j)},s^i,\xi^i)=x^i\right)\\&=P_\xi\left(\xi^i:\Phi^i(\mu^i_\infty,s^i,\xi^i)=\{x^i\}\right),\end{aligned}$$

where the second line follows by optimality of ϕ^i and Lemma 3. This implies that $\sigma^i(x^i|s^i)=P_\xi\left(\xi^i:\Phi^i(\mu^i_\infty,s^i,\xi^i)=\{x^i\}\right)$. Thus, it remains to show that

$$P_\xi\left(\xi^i:\Phi^i(\mu^i_\infty,s^i,\xi^i)=\{x^i\}\right)=P_\xi\left(\xi^i:\psi^i(\mu^i_\infty,s^i,\xi^i)=\{x^i\}\right) \tag{27}$$

$\forall x^i$ such that $P_\xi\left(\xi^i:\Phi^i(\mu^i_\infty,s^i,\xi^i)=\{x^i\}\right)>0$. From now on, fix any such x^i. Since $\sigma^i(x^i\mid s^i)>0$, the assumption that the game is weakly identified implies that $Q^i_{\theta^i_1}(\cdot\mid x^i,s^i)=Q^i_{\theta^i_2}(\cdot\mid x^i,s^i)\ \forall\theta^i_1,\theta^i_2\in\Theta(\sigma)$. The fact that $\mu^i_\infty\in\Delta(\Theta^i(\sigma))$ then implies that

$$B^i(\mu^i_\infty,s^i,x^i,y^i)=\mu^i_\infty \tag{28}$$

$\forall y^i\in\mathbb{Y}^i$. Thus, $\Phi^i(\mu^i_\infty,s^i,\xi^i)=\{x^i\}$ is equivalent to

$$\begin{aligned}&E_{\bar{Q}_{\mu^i_\infty}(\cdot|s^i,x^i)}\left[\pi^i(x^i,Y^i)+\xi^i(x^i)+\delta E_{p_{S^i}}\left[V^i(\mu^i_\infty,S^i)\right]\right]\\&\quad>E_{\bar{Q}_{\mu^i_\infty}(\cdot|s^i,\tilde{x}^i)}\left[\pi^i(\tilde{x}^i,Y^i)+\xi^i(\tilde{x}^i)+\delta E_{p_{S^i}}\left[V^i(B^i(\mu^i_\infty,s^i,\tilde{x}^i,Y^i),S^i)\right]\right]\\&\quad\geq E_{\bar{Q}_{\mu^i_\infty}(\cdot|s^i,\tilde{x}^i)}\left[\pi^i(\tilde{x}^i,Y^i)+\xi^i(\tilde{x}^i)\right]+\delta E_{p_{S^i}}\left[V^i(E_{\bar{Q}_{\mu^i_\infty}(\cdot|s^i,\tilde{x}^i)}\left[B^i(\mu^i_\infty,s^i,\tilde{x}^i,Y^i)\right],S^i)\right]\\&\quad=E_{\bar{Q}_{\mu^i_\infty}(\cdot|s^i,\tilde{x}^i)}\left[\pi^i(\tilde{x}^i,Y^i)+\xi^i(\tilde{x}^i)\right]+\delta E_{p_{S^i}}\left[V^i(\mu^i_\infty,S^i)\right]\end{aligned}$$

$\forall\tilde{x}^i\in\mathbb{X}^i$, where the first line follows by equation (28) and definition of Φ^i, the second line follows by the convexity[50] of V^i as a function of μ^i and Jensen's inequality, and the last line by the fact that Bayesian beliefs have the martingale property. In turn,

[50]See, for example, Nyarko (1994), for a proof of convexity of the value function.

the above expression is equivalent to $\psi(\mu^i_\infty, s^i, \xi^i) = \{x^i\}$. □

D Population models

We discuss some variants of population models that differ in the matching technology and feedback. The right variant of population model will depend on the specific application.[51]

Single pair model. Each period a single pair of players is randomly selected from each of the i populations to play the game. At the end of the period, the signals, actions, and outcomes of their own population are revealed to everyone.[52] Steady-state behavior in this case corresponds exactly to the notion of Berk-Nash equilibrium described in the paper.

Random matching model. Each period, all players are randomly matched and observe only feedback from their own match. We now modify the definition of Berk-Nash equilibrium to account for this random-matching setting. The idea is similar to Fudenberg and Levine's (1993) definition of a heterogeneous self-confirming equilibrium. Now each agent in population i can have different experiences and, hence, have different beliefs and play different strategies in steady state.

For all $i \in I$, define

$$BR^i(\sigma^{-i}) = \left\{\sigma^i : \sigma^i \text{ is optimal given } \mu^i \in \Delta\left(\Theta^i(\sigma^i, \sigma^{-i})\right)\right\}.$$

Note that σ is a Berk-Nash equilibrium if and only if $\sigma^i \in BR^i(\sigma^{-i})\ \forall i \in I$.

Definition 9. A strategy profile σ is a **heterogeneous Berk-Nash equilibrium** of game $\mathcal{G}$ if, for all $i \in I$, σ^i is in the convex hull of $BR^i(\sigma^{-i})$.

Intuitively, a heterogenous equilibrium strategy σ^i is the result of convex combinations of strategies that belong to $BR^i(\sigma^{-i})$; the idea is that each of these strategies

[51] In some cases, it may be unrealistic to assume that players are able to observe the private signals of previous generations, so some of these models might be better suited to cases with public, but not private, information.

[52] Alternatively, we can think of different incarnations of players born every period who are able to observe the history of previous generations.

is followed by a segment of the population i.[53]

Random-matching model with population feedback. Each period all players are randomly matched; at the end of the period, each player in population i observes the signals, actions, and outcomes of their own population. Define

$$\bar{BR}^i(\sigma^i, \sigma^{-i}) = \left\{\hat{\sigma}^i : \hat{\sigma}^i \text{ is optimal given } \mu^i \in \Delta\left(\Theta^i(\sigma^i, \sigma^{-i})\right)\right\}.$$

Definition 10. A strategy profile σ is a **heterogeneous Berk-Nash equilibrium with population feedback** of game $\mathcal{G}$ if, for all $i \in I$, σ^i is in the convex hull of $\bar{BR}^i(\sigma^i, \sigma^{-i})$.

The main difference when players receive population feedback is that their beliefs no longer depend on their own strategies but rather on the aggregate population strategies.

D.1 Equilibrium foundation

Using arguments similar to the ones in the text, it is now straightforward to conclude that the definition of heterogenous Berk-Nash equilibrium captures the steady state of a learning environment with a population of agents in the role of each player. To see the idea, let each population i be composed of a continuum of agents in the unit interval $K \equiv [0, 1]$. A strategy of agent ik (meaning agent $k \in K$ from population i) is denoted by σ^{ik}. The aggregate strategy of population (i.e., player) i is $\sigma^i = \int_K \sigma^{ik} dk$.

Random matching model. Suppose that each agent is optimizing and that, for all i, (σ_t^{ik}) converges to σ^{ik} a.s. in K, so that individual behavior stabilizes.[54] Then Lemma 2 says that the support of beliefs must eventually be $\Theta^i(\sigma^{ik}, \sigma^{-i})$ for agent ik. Next, for each ik, take a convergent subsequence of beliefs μ_t^{ik} and denote it μ_∞^{ik}. It follows that $\mu_\infty^{ik} \in \Delta(\Theta^i(\sigma^{ik}, \sigma^{-i}))$ and, by continuity of behavior in beliefs, σ^{ik} is optimal given μ_∞^{ik}. In particular, $\sigma^{ik} \in BR^i(\sigma^{-i})$ for all ik and, since $\sigma^i = \int_K \sigma^{ik} dk$, it follows that σ^i is in the convex hull of $BR^i(\sigma^{-i})$.

[53] Unlike the case of heterogeneous self-confirming equilibrium, a definition where each action in the support of σ is supported by a (possibly different) belief would not be appropriate here. The reason is that $BR^i(\sigma^{-i})$ might contain only mixed, but not pure strategies (e.g., Example 1).

[54] We need individual behavior to stabilize; it is not enough that it stabilizes in the aggregate. This is natural, for example, if we believe that agents whose behavior is unstable will eventually realize they have a misspecified model.

Random matching model with population feedback. Suppose that each agent is optimizing and that, for all i, $\sigma_t^i = \int_K \sigma_t^{ik} dk$ converges to σ^i. Then Lemma 2 says that the support of beliefs must eventually be $\Theta^i(\sigma^i, \sigma^{-i})$ for any agent in population i. Next, for each ik, take a convergent subsequence of beliefs μ_t^{ik} and denote it μ_∞^{ik}. It follows that $\mu_\infty^{ik} \in \Delta(\Theta^i(\sigma^i, \sigma^{-i}))$ and, by continuity of behavior in beliefs, σ^{ik} is optimal given μ_∞^{ik}. In particular, $\sigma^{ik} \in B\bar{R}^i(\sigma^{-i})$ for all i, k and, since $\sigma^i = \int_K \sigma^{ik} dk$, it follows that σ^i is in the convex hull of $B\bar{R}^i(\sigma^{-i})$.

E Lack of payoff feedback

In the paper, players are assumed to observe their own payoffs. We now provide two alternatives to relax this assumption. In the first alternative, players observe no feedback about payoffs; in the second alternative, players may observe partial feedback.

No payoff feedback. In the paper we had a single, deterministic payoff function $\pi^i : \mathbb{X}^i \times \mathbb{Y}^i \to \mathbb{R}$, which can be represented in vector form as an element $\pi^i \in \mathbb{R}^{\#(\mathbb{X}^i \times \mathbb{Y}^i)}$. We now generalize it to allow for uncertain payoffs. Player i is endowed with a probability distribution $P_{\pi^i} \in \Delta(\mathbb{R}^{\#(\mathbb{X}^i \times \mathbb{Y}^i)})$ over the possible payoff functions. In particular, the random variable π^i is independent of Y^i, and so there is nothing new to learn about payoffs from observing consequences. With random payoff functions, the results extend provided that optimality is defined as follows: A strategy σ^i for player i is **optimal** given $\mu^i \in \Delta(\Theta^i)$ if $\sigma^i(x^i \mid s^i) > 0$ implies that

$$x^i \in \arg\max_{\bar{x}^i \in \mathbb{X}^i} E_{P_{\pi^i}} E_{\bar{Q}^i_{\mu^i}(\cdot \mid s^i, \bar{x}^i)} \left[\pi^i(\bar{x}^i, Y^i) \right].$$

Note that by interchanging the order of integration, this notion of optimality is equivalent to the notion in the paper where the deterministic payoff function is given by $E_{P_{\pi^i}} \pi^i(\cdot, \cdot)$.

Partial payoff feedback. Suppose that player i knows her own consequence function $f^i : \mathbb{X} \times \Omega \to \mathbb{Y}^i$ and that her payoff function is now given by $\pi^i : \mathbb{X} \times \Omega \to \mathbb{R}$. In particular, player i may not observe her own payoff, but observing a consequence may provide partial information about (x^{-i}, ω) and, therefore, about payoffs. Unlike the case in the text where payoffs are observed, a belief $\mu^i \in \Delta(\Theta^i)$ may not uniquely determine expected payoffs. The reason is that the distribution over consequences

implied by μ^i may be consistent with several distributions over $\mathbb{X}^{-i} \times \Omega$; i.e., the distribution over $\mathbb{X}^{-i} \times \Omega$ is only partially identified. Define the set $\mathcal{M}_{\mu^i} \subseteq \Delta(\mathbb{X}^{-i} \times \Omega)^{\mathbb{S}^i \times \mathbb{X}^i}$ to be the set of conditional distributions over $\mathbb{X}^{-i} \times \Omega$ given $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$ that are consistent with belief $\mu^i \in \Delta(\Theta^i)$, i.e., $m \in \mathcal{M}_{\mu^i}$ if and only if $\bar{Q}^i_{\mu^i}(y^i \mid s^i, x^i) = m\left(f^i(x^i, X^{-i}, W) = y^i \mid s^i, x^i\right)$ for all $(s^i, x^i) \in \mathbb{S}^i \times \mathbb{X}^i$ and $y^i \in \mathbb{Y}^i$. Then optimality should be defined as follows: A strategy σ^i for player i is **optimal** given $\mu^i \in \Delta(\Theta^i)$ if there exists $m_{\mu^i} \in \mathcal{M}_{\mu^i}$ such that $\sigma^i(x^i \mid s^i) > 0$ implies that

$$x^i \in \arg\max_{\bar{x}^i \in \mathbb{X}^i} E_{m_{\mu^i}(\cdot \mid s^i, \bar{x}^i)}\left[\pi^i(\bar{x}^i, X^{-i}, W)\right].$$

Finally, the definition of identification would also need to be changed to require not only that there is a unique distribution over consequences that matches the observed data, but also that this unique distribution implies a unique expected utility function.

F Global stability: Example 2.1 (monopoly with unknown demand).

Theorem 3 says that all Berk-Nash equilibria can be approached with probability 1 provided we allow for vanishing optimization mistakes. In this appendix, we illustrate how to use the techniques of stochastic approximation theory to establish stability of equilibria under the assumption that players make no optimization mistakes. We present the explicit learning dynamics for the monopolist with unknown demand, Example 2.1, and show that the unique equilibrium in this example is globally stable. The intuition behind global stability is that switching from the equilibrium strategy to a strategy that puts more weight on a price of 2 changes beliefs in a way that makes the monopoly want to put less weight on a price of 2, and similarly for a deviation to a price of 10.

We first construct a perturbed version of the game. Then we show that the learning problem is characterized by a nonlinear stochastic system of difference equations and employ stochastic approximation methods for studying the asymptotic behavior of such system. Finally, we take the payoff perturbations to zero.

In order to simplify the exposition and thus better illustrate the mechanism driving the dynamics, we modify the subjective model slightly. We assume the monopolist only learns about the parameter $b \in \mathbb{R}$; i.e., her beliefs about parameter a are degenerate at

a point $a = 40 \neq a^0$ and thus are never updated. Therefore, beliefs μ are probability distributions over $\mathbb{R}$, i.e., $\mu \in \Delta(\mathbb{R})$.

Perturbed Game. Let ξ be a real-valued random variable distributed according to P_ξ; we use F to denote the associated cdf and f the pdf. The perturbed payoffs are given by $yx - \xi 1\{x = 10\}$. Thus, given beliefs $\mu \in \Delta(\mathbb{R})$, the probability of optimally playing $x = 10$ is

$$\sigma(\mu) = F\left(8a - 96E_\mu[B]\right).$$

Note that the only aspect of μ that matters for the decision of the monopolist is $E_\mu[B]$. Thus, letting $m = E_\mu[B]$ and slightly abusing notation, we use $\sigma(\mu) = \sigma(m)$ as the optimal strategy.

Bayesian Updating. We now derive the Bayesian updating procedure. We assume that the the prior μ_0 is given by a Gaussian distribution with mean and variance m_0, τ_0^2.[55] It is possible to show that, given a realization (y, x) and a prior $N(m, \tau^2)$, the posterior is also Gaussian and the mean and variance evolve as follows: $m_{t+1} = m_t + \left(\frac{-(Y_{t+1}-a)}{X_{t+1}} - m_t\right)\left(\frac{X_{t+1}^2}{X_{t+1}^2+\tau_t^{-2}}\right)$ and $\tau_{t+1}^2 = \frac{1}{(X_{t+1}^2+\tau_t^{-2})}$.

Nonlinear Stochastic Difference Equations and Stochastic Approximation. For simplicity, let $r_{t+1} \equiv \frac{1}{t+1}\left(\tau_t^{-2} + X_{t+1}^2\right)$ and note that the previous nonlinear system of stochastic difference equations can be written as

$$\begin{aligned} m_{t+1} &= m_t + \frac{1}{t+1}\frac{X_{t+1}^2}{r_{t+1}}\left(\frac{-(Y_{t+1}-a)}{X_{t+1}} - m_t\right) \\ r_{t+1} &= r_t + \frac{1}{t+1}\left(X_{t+1}^2 - r_t\right). \end{aligned}$$

Let $\beta_t = (m_t, r_t)'$, $Z_t = (X_t, Y_t)$,

$$G(\beta_t, z_{t+1}) = \begin{bmatrix} \frac{x_{t+1}^2}{r_{t+1}}\left(\frac{-(y_{t+1}-a)}{x_{t+1}} - m_t\right) \\ \left(x_{t+1}^2 - r_t\right) \end{bmatrix}$$

[55] This choice of prior is standard in Gaussian settings like ours. As shown below this choice simplifies the exposition considerably.

and

$$\mathbb{G}(\beta) = \begin{bmatrix} \mathbb{G}_1(\beta) \\ \mathbb{G}_2(\beta) \end{bmatrix} = E_{P_\sigma}\left[G(\beta, Z_{t+1})\right]$$
$$= \begin{bmatrix} F(8a-96m)\frac{100}{r}\left(\frac{-(a_0-a-b_0 10)}{10} - m\right) + (1-F(8a-96m))\frac{4}{r}\left(\frac{-(a_0-a-b_0 2)}{2} - m\right) \\ (4 + F(8a-96m)96 - r) \end{bmatrix}$$

where P_σ is the probability over Z induced by σ (and $y = a^0 - b^0 x + \omega$). Therefore, the dynamical system can be cast as

$$\beta_{t+1} = \beta_t + \frac{1}{t+1}\mathbb{G}(\beta_t) + \frac{1}{t+1}V_{t+1}$$

with $V_{t+1} = G(\beta_t, Z_{t+1}) - \mathbb{G}(\beta_t)$. Stochastic approximation theory (e.g., Kushner and Yin (2003)) implies, roughly speaking, that in order to study the asymptotic behavior of $(\beta_t)_t$ it is enough to study the behavior of the orbits of the following ODE

$$\dot{\beta}(t) = \mathbb{G}(\beta(t)).$$

Characterization of the Steady States. In order to find the steady states of $(\beta_t)_t$, it is enough to find β^* such that $\mathbb{G}(\beta^*) = 0$. Let $H_1(m) \equiv F(8a - 96m)10\left(-(a_0 - a) + (b_0 - m)\,10\right) + (1 - F(8a - 96m))\,2\left(-(a_0 - a) + (b_0 - m)\,2\right)$. Observe that $\mathbb{G}_1(\beta) = r^{-1}H_1(m)$ and that H_1 is continuous and $\lim_{m\to-\infty} H_1(m) = \infty$ and $\lim_{m\to\infty} H_1(m) = -\infty$. Thus, there exists at least one solution $H_1(m) = 0$. Therefore, there exists at least one β^* such that $\mathbb{G}(\beta^*) = 0$.

Let $\bar{b} = b_0 - \frac{a_0-a}{10} = 4 - \frac{1}{5} = \frac{19}{5}$ and $\underline{b} = b_0 - \frac{a_0-a}{2} = 4 - \frac{42-40}{2} = 3$, $\bar{r} = 4 + F(8a - 96\underline{b})96$ and $\underline{r} = 4 + F(8a - 96\bar{b})96$, and $\mathbb{B} \equiv [\underline{b}, \bar{b}] \times [\underline{r}, \bar{r}]$. It follows that $H_1(m) < 0$ $\forall m > \bar{b}$, and thus m^* must be such that $m^* \leq \bar{b}$. It is also easy to see that $m^* \geq \underline{b}$. Moreover, $\frac{dH_1(m)}{dm} = 96f(8a-96m)\big(8(a_0-a) - \big(b_0-m\big)96\big) - 4 - 96F(8a-96m)$. Thus, for any $m \leq \bar{b}$, $\frac{dH_1(m)}{dm} < 0$, because $m \leq \bar{b}$ implies $8(a_0-a) \leq (b_0-m)80 < (b_0-m)96$.

Therefore, on the relevant domain $m \in [\underline{b}, \bar{b}]$, H_1 is decreasing, thus implying that there exists only one m^* such that $H_1(m^*) = 0$. Therefore, there exists only one β^* such that $\mathbb{G}(\beta^*) = 0$.

We are now interested in characterizing the limit of β^* as the perturbation vanishes, i.e. as F converges to $1\{\xi \geq 0\}$. To do this we introduce some notation. We consider a sequence $(F_n)_n$ that converges to $1\{\xi \geq 0\}$ and use β_n^* to denote the steady state

associated to F_n. Finally, we use H_1^n to denote the H_1 associated to F_n.

We proceed as follows. First note that since $\beta_n^* \in \mathbb{B}\ \forall n$, the limit exists (going to a subsequence if needed). We show that $m^* \equiv \lim_{n\to\infty} m_n^* = \frac{8a}{96} = 8\frac{40}{96} = \frac{10}{3}$. Suppose not, in particular, suppose that $\lim_{n\to\infty} m_n^* < \frac{8a}{96} = \frac{10}{3}$ (the argument for the reverse inequality is analogous and thus omitted). In this case $\lim_{n\to\infty} 8a - 96m_n^* > 0$, and thus $\lim_{n\to\infty} F_n(8a - 96m_n^*) = 1$. Therefore

$$\lim_{n\to\infty} H_1^n(\beta_n^*) = 10\left(-(a_0 - a) + (b_0 - m^*)\,10\right) \geq 10\left(-2 + \frac{2}{3}10\right) > 0.$$

But this implies that $\exists N$ such that $H_1^n(\beta_n^*) > 0\ \forall n \geq N$ which is a contradiction.

Moreover, define $\sigma_n^* = F_n(8a - 96m_n^*)$ and $\sigma^* = \lim_{n\to\infty} \sigma_n$. Since $H_1^n(m_n^*) = 0\ \forall n$ and $m^* = \frac{10}{3}$, it follows that

$$\sigma^* = \frac{-2\left(-2 + \left(4 - \frac{10}{3}\right)2\right)}{10\left(-2 + \left(4 - \frac{10}{3}\right)10\right) - 2\left(-2 + \left(4 - \frac{10}{3}\right)2\right)} = \frac{1}{36}.$$

Global convergence to the Steady State. In our example, it is in fact possible to establish that behavior converges with probability 1 to the unique equilibrium. By the results in Benaim (1999) Section 6.3, it is sufficient to establish the *global* asymptotic stability of β_n^* for any n, i.e., the basin of attraction of β_n^* is all of $\mathbb{B}$.

In order to do this let $L(\beta) = (\beta - \beta_n^*)' P (\beta - \beta_n^*)$ for all β; where $P \in \mathbb{R}^{2\times 2}$ is positive definite and *diagonal* and will be determined later. Note that $L(\beta) = 0$ iff $\beta = \beta_n^*$. Also

$$\begin{aligned}
\frac{dL(\beta(t))}{dt} &= \nabla L(\beta(t))' \mathbb{G}(\beta(t)) \\
&= 2\left(\beta(t) - \beta_n^*\right)' P \left(\mathbb{G}(\beta(t))\right) \\
&= 2\left\{(m(t) - m_n^*)\, P_{[11]} \mathbb{G}_1(\beta(t)) + (r(t) - r_n^*)\, P_{[22]} \mathbb{G}_2(\beta(t))\right\}.
\end{aligned}$$

Since $\mathbb{G}(\beta_n^*) = 0$,

$$\begin{aligned}
\frac{dL(\beta(t))}{dt} =& 2\left(\beta(t) - \beta_n^*\right)' P\left(\mathbb{G}(\beta(t)) - \mathbb{G}\left(\beta_n^*\right)\right) \\
=& 2\left(m(t) - m_n^*\right) P_{[11]}\left(\mathbb{G}_1(\beta(t)) - \mathbb{G}_1\left(\beta_n^*\right)\right) \\
&+ 2\left(r(t) - r_n^*\right) P_{[22]}\left(\mathbb{G}_2(\beta(t)) - \mathbb{G}_2\left(\beta_n^*\right)\right) \\
=& 2\left(m(t) - m_n^*\right)^2 P_{[11]} \int_0^1 \frac{\partial \mathbb{G}_1(m_n^* + s(m(t) - m_n^*), r_n^*)}{\partial m} ds \\
&+ 2\left(r(t) - r_n^*\right)^2 P_{[22]} \int_0^1 \frac{\partial \mathbb{G}_2(m_n^*, r_n^* + s(r(t) - r_n^*))}{\partial r} ds
\end{aligned}$$

where the last equality holds by the mean value theorem. Note that $\frac{d\mathbb{G}_2(m_n^*, r_n^* + s(r(t) - r_n^*))}{dr} = -1$ and $\int_0^1 \frac{d\mathbb{G}_1(m_n^* + s(m(t) - m_n^*), r_n^*)}{dm} ds = \int_0^1 (r_n^*)^{-1} \frac{dH_1(m_n^* + s(m(t) - m_n^*))}{dm} ds$. Since $r(t) > 0$ and $r_n^* \geq 0$ the first term is positive and we already established that $\frac{dH_1(m)}{dm} < 0 \; \forall m$ in the relevant domain. Thus, by choosing $P_{[11]} > 0$ and $P_{[22]} > 0$ it follows that $\frac{dL(\beta(t))}{dt} < 0$.

Therefore, we show that L satisfies the following properties: is strictly positive $\forall \beta \neq \beta_n^*$ and $L(\beta_n^*) = 0$, and $\frac{dL(\beta(t))}{dt} < 0$. Thus, the function satisfies all the conditions of a Lyapounov function and, therefore, β_n^* is globally asymptotically stable $\forall n$ (see Hirsch et al. (2004) p. 194).